

High-Accuracy Laser Power and Energy Meter Calibration Service

by

David Livigni

Contents

1. INTRODUCTION	1
2. HIGH-ACCURACY LASER POWER AND ENERGY METER CALIBRATION SYSTEM	2
2.1 Calibration System Overview	2
2.2 Calibration Philosophy	4
2.3 Parameter Specification for End-User Uncertainty Assessment	5
3. SYSTEM DESIGN AND IMPLEMENTATION	8
3.1 High-Accuracy Calibration System	8
3.1.1 Data acquisition system	8
3.1.2 Laboratory environment	8
3.1.3 Mechanical systems	10
3.1.4 Gas and cryogen handling systems	12
3.1.5 Optical system	12
3.1.6 Laser power stabilization system	16
3.1.7 The relative aperture transmittance correction	18
3.2 The LOCR Primary Standard	19
3.2.1 Basic cryogenic radiometer operating principal	19
3.2.2 Description of the LOCR	20
3.2.3 LOCR electrical calibration	23
3.2.4 The NIST Brewster's angle window	25
4. CALIBRATION PROCEDURE	28
4.1 Calibration System Operation	28
4.1.1 Data acquisition software	28
4.1.2 Optical component selection and placement	28
4.1.3 Laser source alignment	29
4.1.4 Detector alignment	31
4.1.5 Relative aperture transmittance measurement	33
4.1.6 Calibration system maintenance	35
4.1.7 Quality assurance checks	36
4.2 LOCR Operation	37
4.2.1 Vacuum and cryogenic systems	37
4.2.2 Alignment and spatial uniformity measurement	42
4.2.3 Electrical calibration	44
4.2.4 Window transmittance measurement	45
4.2.5 Data acquisition	46
4.2.6 LOCR maintenance	48
4.3 Calibration Factor Calculation	49
4.3.1 Meter output analysis	50
4.3.2 Applied power calculation	51
4.3.3 Calibration factor calculation	52
5. UNCERTAINTY ANALYSIS	54
5.1 Uncertainty of the LOCR Primary Standard	55
5.1.1 Window transmittance	56
5.1.2 Receiver absorptance	56

5.1.3 Receiver alignment	57
5.1.4 Electrical calibration	57
5.1.5 Electrical heating inequivalence	57
5.1.6 Other LOCR uncertainties	58
5.2 Uncertainty of the Applied Power	59
5.2.1 Relative aperture transmittance	59
5.2.2 Measurement repeatability of the applied power	59
5.2.3 Uncertainty of the applied energy	60
5.3 Uncertainty of the Test Detector Calibration	60
5.3.1 NIST amplifier gain	60
5.3.2 NIST electrical measurement	61
5.3.3 Quantization in digital test meters	62
5.3.4 Measurement repeatability of the calibration factor	64
5.3.5 Uncertainty in the calibration factor for energy meters	64
5.4 Other Uncertainties	64
5.4.1 Test detector alignment	64
5.4.2 Laser beam spatial characteristics	65
5.4.3 Laser beam wavelength and spectrum	67
5.4.4 Environment specification	68
REFERENCES	70
APPENDIX A. Glossary	72
APPENDIX B. Instructions for Submitting Meters for Calibration	76
APPENDIX C. Sample Calibration Report	78
APPENDIX D. Gaussian Beam Model and Spatial Filter Pinhole Selection	86
D.1 Geometrical Optics Analysis	86
D.2 Gaussian Beam Ray-Tracing Analysis	87
D.3 Gaussian Beam Parameter Specification	95
D.4 Spatial Filter Pinhole Selection	96
APPENDIX E. Relative Aperture Transmittance: Estimation and Uncertainty	99
E.1 Measured and Theoretical Transmittance for Circular Apertures	100
E.2 Relative Transmittance Measurement for Circular Apertures	104
E.3 Estimation of the Relative Aperture Transmittance for Detectors	107
APPENDIX F. System DMM and LOCR Electronics: Calibration and Correction	111
F.1 System DMM Calibration	111
F.2 Measurement of LOCR Electrical Power	113
F.3 Calibration of LOCR Power	115
F.4 Combined LOCR Power Measurement Correction	118
APPENDIX G. LOCR Optical Receiver Alignment: Uncertainty	119
G.1 Philosophy of Receiver Uniformity Measurement	119
G.2 Receiver Uniformity Measurement: Technique and Analysis	120
APPENDIX H. Window Transmittance: Estimation and Uncertainty	123

H.1 Window Transmittance Measurement: Philosophy and Technique	123
H.2 Transmittance Measurement: Uncertainty Analysis	125
H.3 Example Calculation of Transmittance	128
APPENDIX I. Center-Weighted Spatial Uncertainty	129
APPENDIX J. Estimate of LOCR Optical Receiver Absorptance	132
J.1 LOCR Receiver Absorptance at 633 nm	132
J.2 LOCR Receiver Absorptance at Other Wavelengths	133
APPENDIX K. Weather Station Calibration and Uncertainty	136
K.1 Internal and External Temperature Sensor Calibration	136
K.2 Barometric Pressure Sensor Calibration	139
K.3 Relative Humidity Sensor Calibration	142

List of Illustrations

Figure 1. Calibration system overview.	2
Figure 2. Current configuration of the high-accuracy calibration system.	7
Figure 3. High-accuracy calibration system laboratory table and enclosure.	9
Figure 4. High-accuracy calibration system's laboratory table layout.	11
Figure 5. Detailed view of the optical source breadboard.	13
Figure 6. The apparent aperture of a detector that has no limiting aperture.	19
Figure 7. External view of the LOCR.	20
Figure 8. Internal view of the LOCR.	21
Figure 9. The LOCR's optical receiver.	22
Figure 10. General functional schematic of the LOCR.	22
Figure 11. Simplified LOCR electrical calibration circuit.	23
Figure 12. The NIST Brewster's angle window mount.	25
Figure 13. Main calibration process flow chart, for calibration at a single wavelength.	27
Figure 14. Rotational and translational motions used during detector alignment.	32
Figure 15. The NIST Brewster's angle window mount, showing the vertical offset of the incident laser beam.	43
Figure 16. Typical LOCR illuminated output signal. The signal is dithered, so the quantization error is randomized.	58
Figure 17. The output current from an illuminated photodiode, showing two discrete quantization levels.	63
Figure D.1 Thin lens Gaussian optics model showing the 1550 nm optical source; mirrors, polarizers, power modulator elements, and aperture stops are not shown.	86
Figure D.2 The size (radius) of the Gaussian beam in the detector plane versus the displacement of Lens 2 from the optimally collimated location.	92
Figure D.3 The radius of curvature of the wave front in the detector plane versus the displacement of Lens 2 from the optimally collimated location.	93
Figure D.4 The distance from Lens 2 to the projected waist image (l_2') versus the displacement of Lens 2 from the optimally collimated location.	94
Figure D.5 The radius of curvature of the wave front in the detector plane versus the displacement in Lens 2 from the optimally collimated location, when the 1550 nm laser is moved 1 m farther away from Lens 1.	94
Figure E.1 Measured and theoretical relative aperture transmittance for small circular apertures, relative to a circular reference aperture of 10 mm diameter, using a beam of 2 mm diameter.	101
Figure E.2 Measured and theoretical relative aperture transmittance for large circular apertures, relative to a circular reference aperture of 10 mm diameter, using a beam of 2 mm diameter.	101
Figure K.1 Internal weather station temperature during a typical optical calibration.	138
Figure K.2 External weather station temperature during a typical optical calibration.	138
Figure K.3 The indicated barometric pressure acquired during a typical optical calibration.	141
Figure K.4 The relative humidity measured during a typical optical calibration.	143

List of Acronyms

CO ₂	Carbon Dioxide
CW	Continuous Wave
DMM	Digital Multimeter
DUT	Device Under Test
He	Gaseous Helium
HEPA	High Efficiency Particulate Air
LOCR	Laser Optimized Cryogenic Radiometer
LPC	Laser Power Controller
N ₂	Gaseous Nitrogen
O ₂	Gaseous Oxygen
OSA	Optical Spectrum Analyzer
ppm	Parts per million
UPS	Uninterruptible Power Supply

High-Accuracy Laser Power and Energy Meter Calibration Service

David J. Livigni
National Institute of Standards and Technology
Boulder, CO 80303

This document describes the high-accuracy laser power and energy meter calibration service provided by the National Institute of Standards and Technology (NIST). Calibrations are performed by direct substitution of a test detector with a cryogenic laser radiometer, traceable to NIST electrical standards. The service currently supports measurements with laser powers from 0.1 to 1.0 mW, at several vacuum wavelengths in the range from 458 to 1550 nm. The expanded uncertainty (with a coverage factor of $k = 2$) of calibrations based on the primary standard typically ranges from 0.02 to 0.05 %. A detailed description of the measurement system and uncertainty analysis is presented.

KEY WORDS: absolute power, calibration, cryogenic radiometer, laser power, optical power.

1. INTRODUCTION

We have built a calibration system to meet the industry's need for very high-accuracy measurements of laser power and energy, at various (selected) wavelengths. The system employs a commercial cryogenic radiometer as the primary standard. This commercial radiometer is referred to as the Laser Optimized Cryogenic Radiometer (LOCR) and is one of several NIST cryogenic radiometers that are used as primary standards for various types of high-accuracy optical measurements. The calibration system provides measurements of detector responsivity as a function of laser power with an expanded uncertainty of typically 0.03 % or less (with a coverage factor of $k = 2$) and provides traceability to SI units with reduced uncertainty for customers requiring the highest level of accuracy. The new calibration system and the procedures for detector calibration based on the LOCR are described within this document. Some of the terminology used in this document is not commonly used, so the definitions are included in Appendix A.

Optical power and energy measurements are traditionally tied to SI units through electrical standards. This is accomplished by means of optical radiometers and calorimeters, designed to allow accurate comparison of absorbed optical power with dissipated electrical power. Electrical power is applied to an optical receiver as a means to calibrate the induced temperature rise of the receiver as a function of dissipated electrical power. Room-temperature devices of this type have a combined standard uncertainty that is typically limited to a few tenths of a percent. The development of electrically calibrated cryogenic radiometers operating near liquid-helium temperatures [1-3] has led to a more than tenfold increase in the level of accuracy for optical power measurements. Indeed, most industrialized countries have at least one cryogenic radiometer in their national standards laboratories.

Since 1967, NIST has built and maintained room-temperature, electrically calibrated laser calorimeters for the calibration of laser power and energy meters for customers. A number of standard laser calorimeters have been developed to provide measurements over a wide range of laser wavelengths (ultraviolet to far-infrared), power levels (nanowatts to kilowatts), and energy levels (femtojoules to megajoules). The combined standard uncertainty of measurements with these calorimeters is usually limited to about 0.25 %, primarily due to their operation at room temperature. Their performance is limited primarily by inequivalence between electrical and optical heating, due to factors such as: radiative and convective cooling of the optical receiver, and its limited diffusivity, which result in the formation of temperature gradients, and parasitic heating in the electrical heater leads. Commercial laser power and energy meters have significantly improved over the last 10 years and customers now require lower uncertainties. The LOCR calibration service was developed by NIST's Optoelectronics Division to meet select customers' needs for higher accuracy.

2. HIGH-ACCURACY LASER POWER AND ENERGY METER CALIBRATION SYSTEM

The goal of the high-accuracy calibration service is to provide laser power and energy meter calibrations for the most demanding customers, with low uncertainty ($< 0.1\%$), performed in a stable and fully documented environment, at an affordable cost and in a timely manner. The calibration system achieves the goal by offering calibrations based on the highly accurate laser optimized cryogenic radiometer (LOCR) primary standard, using a stabilized laser source, in an automated calibration system that allows quick calibration of multiple detectors at the same time. The service will also be used internally by NIST to improve the accuracy of routine calibrations.

2.1 Calibration System Overview

High-accuracy calibrations are possible because of the accuracy of the primary standard's laser power measurement and the power-stabilized laser source which allows for accurate transfer of the power measurement to the customer's detector. For brevity in the following discussion, the device being calibrated, whether a detector head or complete power or energy meter, is called the device under test (DUT). An overview of the calibration system is shown in Figure 1. The calibration system consists of a power-stabilized laser source, two movable platforms, and a computer-controlled data-acquisition system. The laser source contains a polarized laser, interference filters to remove unwanted laser lines (when necessary), a power stabilization system and an optical spatial filter. The DUT is calibrated using a direct substitution method-the laboratory standard is physically substituted with the DUT. This method requires a stable laser source power, so the power stabilization system is necessary.

The detectors are mounted on a motion-controlled platform that translates in a direction transverse to the laser beam's propagation axis, so that each detector can in turn be moved into the beam's path, thereby

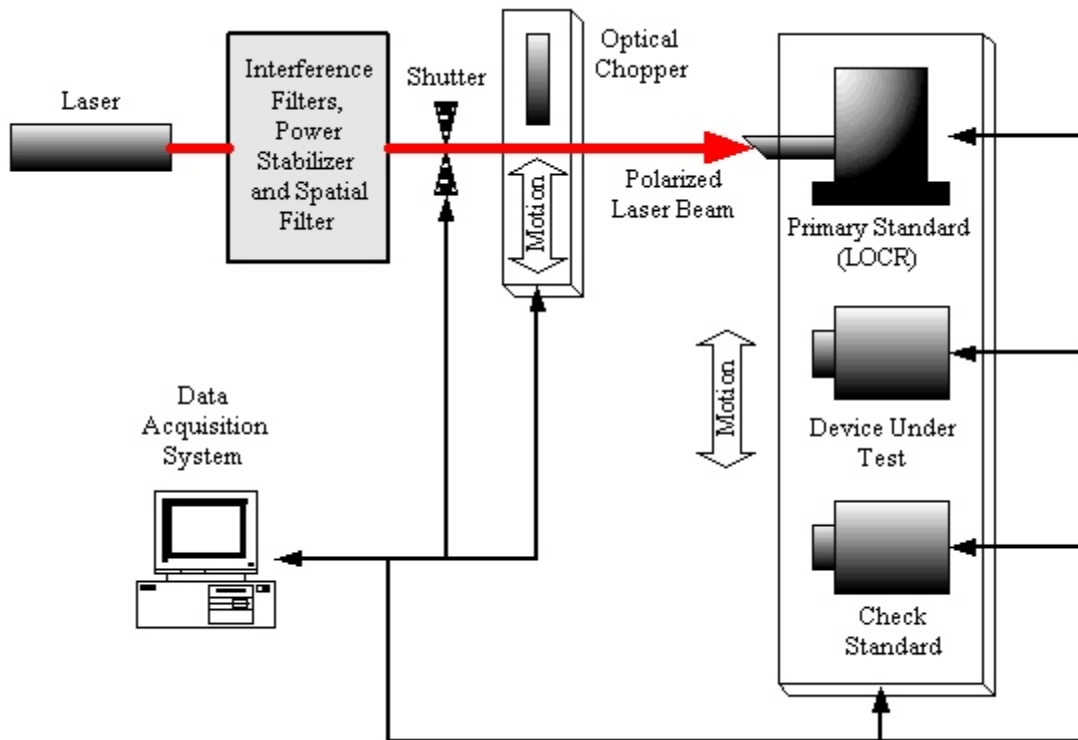


Figure 1. Calibration system overview.

providing the direct substitution of the standard with the DUT. The LOCR primary standard and multiple DUTs can be placed on the platform, allowing calibration of more than one DUT at a time. Each detector is aligned so that when it is moved into the beam, its entrance or limiting aperture is at a specific location along the beam's longitudinal axis, so that each detector is exposed to the same beam profile. A trusted check standard can be calibrated along with the DUT, as a check on quality control for the system.

The LOCR uses an optical receiver design that optimizes the capture of the collimated laser beam and uses a Brewster's angle window that reflects very little of the incident laser power. Polarized laser light passes through the Brewster's angle window with minimal attenuation and is absorbed in the optical receiver, which converts the optical power to heat. The amount of heating power is then measured by electrical substitution, so the device is a primary standard traceable to SI units through the electrical substitution. The laser power that is applied to the DUT is interpolated in time from bracketing measurements performed with the LOCR. The linear interpolation method is used to eliminate the effect of linear drift in the applied optical power.

Energy calibrations are essentially the same as power calibrations, except that the period during which the power is applied to the DUT is also measured. A digital counter records the period over which the optical shutter is open. The applied energy is then calculated by multiplying the interpolated power by the period during which the power is applied. The corrected rise calculation [4], used with many laser calorimeters to measure energy, is also supported.

The calibration system's laser source provides a polarized laser beam with stable continuous-wave (CW) power at discrete laser wavelengths, but an optical chopper can be used with DUTs that require a chopped signal, as shown in Figure 1. The beam's absolute power is easily adjusted using the power stabilizer's electronics and external attenuators, and by adjusting the laser's output power (when possible). When operating at a power level of 1 mW, the stabilized source provides laser power with a standard deviation of typically 0.001 %. Currently, 0.1 mW is the lowest power supported; but even lower powers can be used, but with increasing uncertainty. The beam's polarization axis is fixed in the vertical direction, but the angle can be adjusted slightly by rotating the system's polarizer.

The laser source produces a good quality Gaussian beam, the diameter of which is easily adjusted by adjusting the optical system's lenses. But, because of distortion, artifacts, and scattering produced by the imperfect optical elements and finite apertures that are present in the system, the actual beams used are only approximately Gaussian beams. The uncertainties introduced by the imperfect Gaussian beam are evaluated and accounted for in the high-accuracy calibration system.

The range of laser wavelengths, approximately 425 to 1700 nm, the system can support is limited by the power stabilizer and other optical elements. Currently we can provide lasers with vacuum wavelengths of 458.06, 476.62, 488.13, 496.65, 514.67, 632.99, 1064.42, 1343.09, and 1550.43 nm. New laser wavelengths can be added, as required by customer demand. Support for ultraviolet and far-infrared wavelengths is possible with additional optical equipment.

DUTs that output an analog voltage, have a GPIB interface, or have a direct-reading output are currently supported. Support for devices with an RS-232 serial output will also be added in the next version of the system. The calibration system can also measure sensor devices during the calibration. Examples of sensor measurements include monitoring detector bias voltage and precision temperature measurements. Sensors and DUTs that produce an analog voltage are measured using the system's digital multimeter (DMM); GPIB devices are read directly by the data acquisition computer, but the operator must manually enter data acquired from direct-reading devices. Also, an autonomous weather station is located in the

laboratory. The weather station logs air temperature, pressure, and humidity, so no additional sensors are needed for these environmental parameters.

Calibrations at multiple power levels can be used to determine the DUT's power linearity with a high degree of accuracy, within the limited power range supported by the primary standard. Alternatively, the customer can combine the absolute power with the DUT's known power linearity, to adjust the calibration factor for power levels other than the level used in the calibration. Similarly, calibrations at multiple wavelengths can be used to determine the DUT's spectral responsivity over a limited range.

2.2 Calibration Philosophy

Laser power and energy meters are calibrated without specific knowledge of many of the DUT's parameters. However, evaluation of many of the uncertainties associated with the DUT's calibration does require specific knowledge of some of the DUT's parameters. Some examples are the DUT's spectral responsivity, power linearity, temperature coefficient, spatial uniformity, and sensitivity to beam parameters such as polarization, incidence and divergence angles. Similarly, the uncertainty in the DUT's gain is unknown. To avoid having to gather the information from the customer or measure the DUT's properties ourselves, we specify in the calibration report the conditions present at the time of the calibration.

Because of the DUT's potential sensitivities to these parameters, the customer should provide more detailed information than usual when submitting a detector for high-accuracy calibration. The customer should specify what parameters they desire, such as beam parameters including diameter, divergence, polarization angle, and power. Also important are parameters such as detector orientation, alignment technique, cleanliness, electrical setup, and output signal processing method. The customer should also specify what cleaning technique, if any, should be used with the detector; using an improper cleaning technique can damage the detector or otherwise change its properties. Specific instructions for submitting a detector for high-accuracy calibration are given in Appendix B.

There are two exceptions to this rule, in which the DUT's uncertainty results from a property of the calibration system that interacts with a known property of the DUT. The exceptions are for the correction for the relative aperture transmittance correction and the DUT's quantization uncertainty. The relative aperture transmittance correction results from imperfections in the calibration system's near-Gaussian beam interacting with the fixed, finite entrance apertures of the standard and DUT. Quantization uncertainty is caused by the calibration beam's power varying less than the DUT's quantization interval.

To calculate the correction for relative aperture transmittance, knowledge of the calibration system's actual beam profile and the standard and DUT's apparent apertures is required. The correction is applied by NIST because it is much easier to measure and correct the relative transmittance at the time of calibration, than to provide sufficient information about the calibration beam's profile and the standard's aperture to the customer.

Quantization is an issue with the digital electronics used in some DUTs. Some devices return a limited number of digits in their power readings, which are insufficiently dithered. For such meters, the stability of the calibration system's beam power can result in multiple power measurements that have a standard deviation of zero. In this case, the actual power can be anywhere within the DUT's quantization interval, so an uncertainty for its quantization error is assessed. For digital DUT's in which the readings are dithered, either internally or because of sufficient noise in the applied power, the quantization uncertainty is reduced by averaging and becomes part of the measurement reproducibility. The necessity of assessing

a separate uncertainty for DUT quantization is determined by visual analysis of the readings acquired during the DUT's calibration.

2.3 Parameter Specification for End-User Uncertainty Assessment

Customers who desire high-accuracy calibrations typically use the DUT as a transfer standard in their own laboratory to calibrate other detectors. The customer's calibration system introduces uncertainty into the calibrations they perform with the NIST-calibrated transfer standard. Given the low uncertainty in the high-accuracy calibration provided by NIST, small sources of uncertainty that are usually negligible in other calibrations can become significant and so must be accounted for. To aid the customer in assessing these uncertainties, we specify parameters that are not usually provided with lower-accuracy calibrations.

An example high-accuracy calibration report is given in Appendix C. The parameters specified include a description of the laser beam used in the calibration, details of the detector alignment, and the environmental parameters measured by the autonomous weather station during the calibration. Some quantities are derived, such as ideal Gaussian beam parameters for the applied laser beam and the local air's index of refraction.

The diameter of the laser beam used in the NIST calibration is important, because a change in the beam size used with the detector may change its calibration factor. The beam diameter at the detector's entrance aperture, measured from its $1/e^2$ intensity points, is adjustable from about 0.5 to 4 mm. The diameter is measured using a scanning-slit technique. The diameter of the beam on the detector's actual photosensitive element or elements can be different, depending on how far behind the entrance aperture the elements are and how much the beam diverges. To provide the customer with enough information to calculate how the beam changes with distance, we specify the derived full-width beam divergence angle at the detector's entrance aperture and the parameters necessary for modeling the approximately Gaussian beam.

The uncertainties in the centering of the beam in the detector's entrance aperture and in the angle of incidence of the beam entering the detector are specified because they can impact where the beam falls on the detector's photosensitive element or elements. Uncertainty in the location of the beam in the detector's aperture can cause uncertainty in the detector's responsivity, because of imperfections in the detector's spatial uniformity. The angle of incidence can cause increased uncertainty in the location of the beam, depending on how far behind the entrance aperture the detector element or elements are, and some detectors are sensitive to the angle of incidence of the light itself.

The orientation of the detector relative to the beam's polarization angle is also specified, because some detectors have a different responsivity for different polarizations. The polarization angle in the NIST calibration system is approximately vertical relative to the optical bench. Statements such as "the detector was mounted with its mounting hole down" and "the beam was vertically polarized" describe the orientation. This information tells the customer the basic orientation of the DUT with regard to the beam polarization, but if the DUT's responsivity is sensitive to the beam's polarization angle, additional uncertainty analysis and a more precise measurement of the polarization angle relative to the detector may be necessary.

Laser wavelength and spectrum are specified, because some DUTs are spectrally sensitive. We use single-longitudinal-mode lasers when possible, because they produce a very narrow spectral line, with a center wavelength that can be measured very precisely. But multiple-longitudinal-mode lasers may also be used because of their greater availability and lower cost. An upper bound for the half-power linewidth of the main spectral line is stated, along with the center wavelength. The center wavelength of lasers with

a single dominant narrow spectral line (narrower than 15 THz, about 10 nm at 633 nm) is measured with a precision wavelength meter. Some lasers produce multiple spectral lines; their spectra are usually described in the text by stating the distance between the different lines and their relative intensities. When the main spectral line is too broad or other lines with sufficient intensity are present, the precision wavelength meter cannot be used. The wavelength of such lasers is measured using an optical spectrum analyzer, with a corresponding increase in uncertainty. If a laser's spectrum is too complicated to be described succinctly, a graph of the spectrum as measured with an optical spectrum analyzer can be included in the report.

The laser's wavelength in vacuum is specified, since it is invariant as opposed to the air wavelength, which changes with the air's index of refraction. Also, spectrally sensitive detectors that correct for wavelength internally usually require the vacuum wavelength. The provided local air index can be used to derive the air wavelength from the vacuum wavelength, if desired. Currently five lasers are supported, a 633 nm He-Ne laser, 1064 nm, 1340 nm, and 1550 nm solid-state lasers, and a visible-wavelength tunable argon-ion laser.

The temperature, barometric pressure, and relative humidity present during the NIST calibration are measured using the system's autonomous weather station, and are specified in the calibration report. Temperature usually has the greatest effect on the responsivity of most detectors. Therefore, if the detector is used at a temperature different from that during the NIST calibration, an adjustment to its calibration factor may be necessary, depending on the magnitude of the temperature difference, the DUT's temperature coefficient, and the customer's uncertainty tolerance. Also some detectors, such as the LOCR, are sensitive to temperature change, so the room temperature is held constant during the calibration. The weather station measures the air temperature near the detectors, but a precision analog temperature sensor can be placed in direct contact with the detector for a more precise measurement, if desired. Similarly, a change in the relative humidity or pressure can also affect a detector's responsivity. For example, high humidity can cause problems with detectors; such problems range from a slight change in the detector's absorptance, to condensation forming on cooled parts. The humidity at the NIST laboratory in Boulder, Colorado is not controlled. The relative humidity in the laboratory is usually between 20 and 50 %, but occasionally rises to higher levels.

The index of refraction of the air at the time of the calibration is calculated [5] and specified in the calibration report. A change in the air's refractive index can affect the performance of some DUTs, usually by changing the amount of light that is reflected from the detector. While local changes in the refractive air index are usually very small, the index can be significantly different at other locations. Since the Boulder branch of NIST is located at a relatively high altitude, the difference from that at sea level is significant. The refractive index of air in Boulder is typically 1.00022. The index is calculated using the measured environmental parameters, assuming typical values for other parameters, such as CO₂ content. The uncertainty in the calculated index is dominated by the uncertainty in the measured environmental parameters.

Additional parameters are specified for particular DUTs. For example, some detectors require application of a bias voltage using an external power supply, and the magnitude of the applied voltage can affect the detector's responsivity. So if an external bias voltage is applied by NIST, the voltage is measured and stated in the calibration report. Some detectors require a chopped signal. For these detectors, the chopper frequency can be measured and specified. Similarly, when detectors require unusual tuning or alignment, the procedures used to tune or align the device are also specified. If a NIST amplifier or measurement device is used with the customer's detector, the uncertainty of the NIST equipment is assessed and specified. The uncertainty of customer-supplied equipment is not assessed.

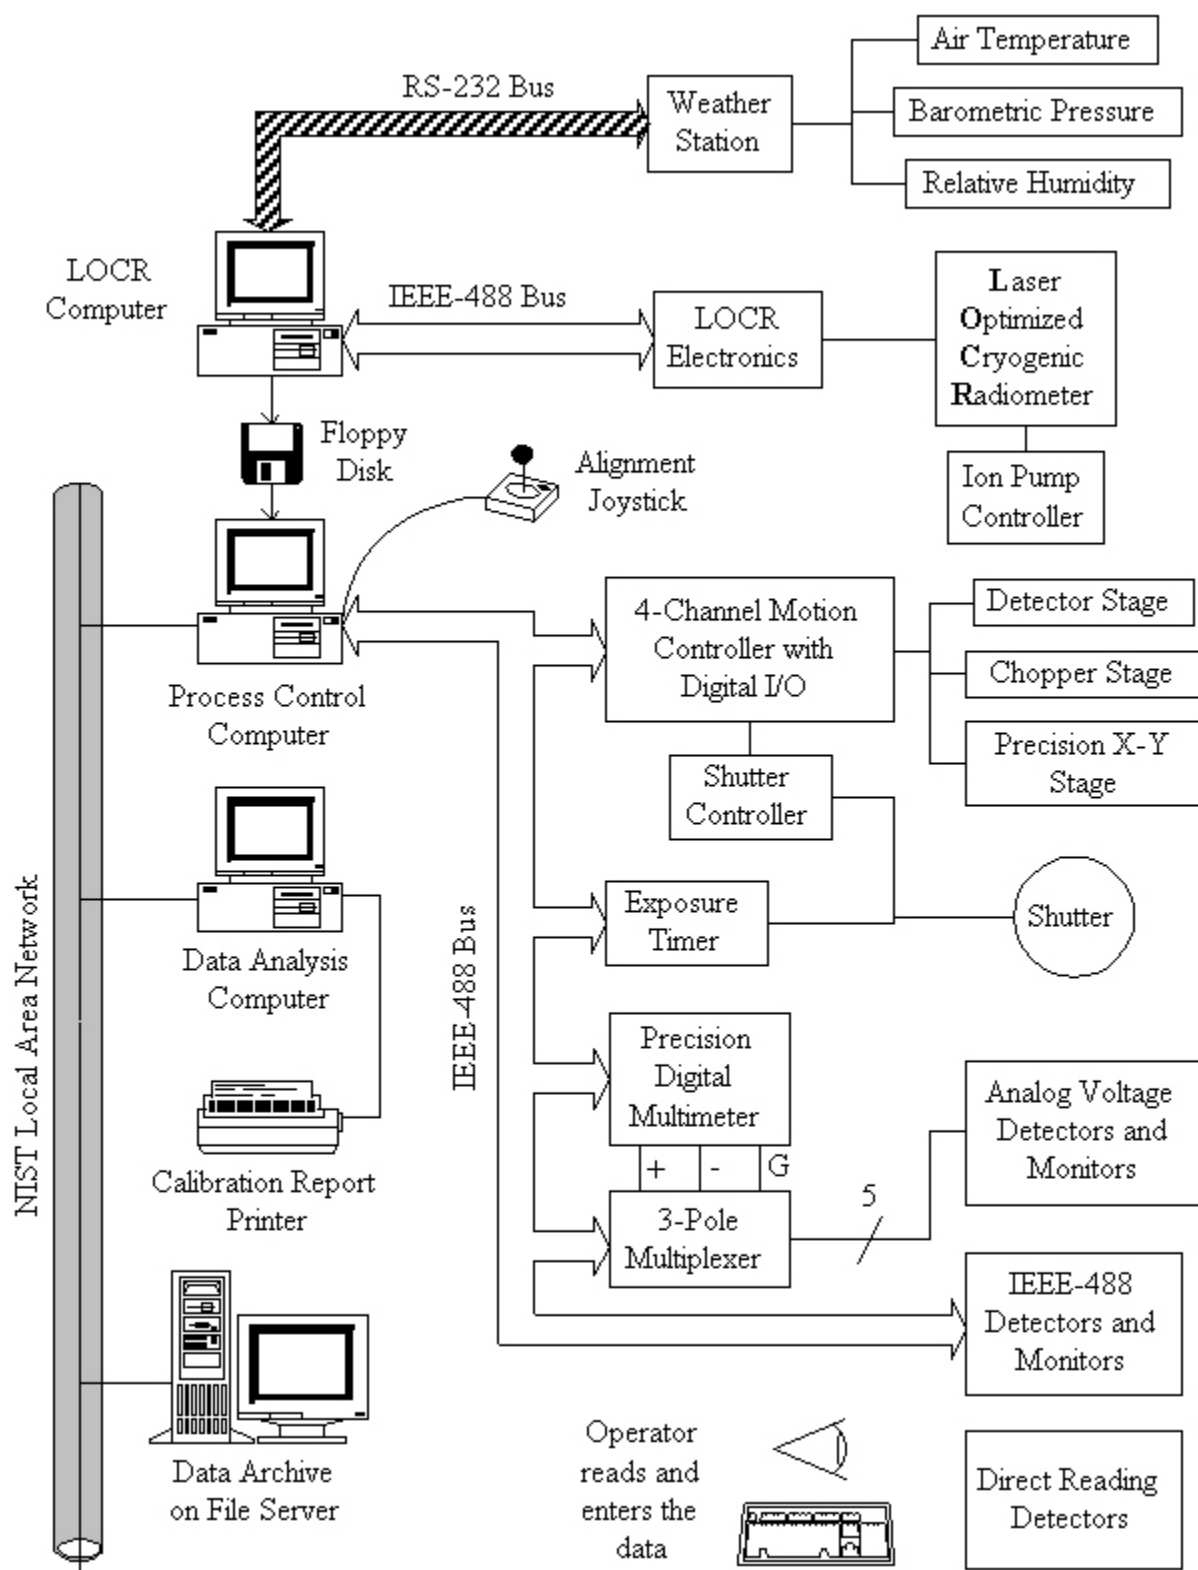


Figure 2. Current configuration of the high-accuracy calibration system.

3. SYSTEM DESIGN AND IMPLEMENTATION

While conceptually simple, performing calibrations with the lowest possible uncertainty is difficult. Many sources of small uncertainties, typically considered negligible in less accurate calibration systems, are not negligible in the high-accuracy calibration system. The high-accuracy system was designed to minimize and account for the uncertainty caused by external influences.

3.1 High-Accuracy Calibration System

A block diagram of the high-accuracy calibration system is shown in Figure 2. Development of the system was frozen in this configuration for introduction to the measurement community at the NewRad 97 conference [6] and so that special test calibrations could be performed for customers.

3.1.1 Data acquisition system: The data acquisition system is controlled by the process control computer, which is programmed using Visual Basic. It controls all the motion control hardware, timing, acquisition, and storage of the test meter power readings. Test meters with analog voltage, digital GPIB, and direct-reading outputs are supported. The system can also acquire data from sensor devices, such as temperature probes or power supply voltages. The acquired data are stored in spreadsheet format and the calculations are performed on the spreadsheet, so that the data processing is fully documented. The LOCR itself is treated as a direct-reading meter that is manually operated using the separate LOCR computer.

3.1.2 Laboratory environment: The laboratory was designed to minimize environmental influences on the calibration system. A precision room-temperature control system is used in the lab, which regulates the room temperature to within a few tenths of a degree Celsius. An autonomous commercial weather station is used to monitor and document the laboratory environment. The weather station measures and records the air temperature, relative humidity, and barometric pressure using internal sensors. The weather station also has an external air temperature sensor that is placed near the detectors to record the air temperature near them. The weather station itself resides on the lab table, where it measures the conditions near the laser source. A significant temperature gradient can exist between the internal and external temperature sensors, depending in part on how much heat is produced by the laser on the lab table. The weather station records the measured environmental parameters internally; the recorded data is periodically transferred to the process control computer, at least once a day. The weather station's temperature sensors have a resolution of about 0.05°C . When necessary, greater temperature precision can be obtained using stand-alone temperature sensors, monitored directly by the calibration system.

To keep dirt and dust off the lab table and optics, the table is enclosed and a High Efficiency Particulate Air (HEPA) filter is used to force clean air into the enclosure. The lab table enclosure is shown in Figure 3; it provides a clean and dark optical environment for the detectors. One obvious sign of the enclosure's effectiveness is a lack of dust particles, which are commonly observed passing through visible laser beams; such dust can add to the laser power instability as the particles drift through the beam. The air flow through the enclosure is not laminar because eddies are formed inside the enclosure, but the improvement in the cleanliness of the lab table is significant. The HEPA filter housing also contains fluorescent lights for illumination of the lab table during setup and an air ionizer that prevents the buildup of static charge.

A problem with static charge was noted when the few dust particles inside the enclosure were found to stick to the optical elements. The particles can become charged and can cling to metal and glass. They can be difficult to remove; rubbing with a brush treated with radioactive material is partially effective, but




Figure 3. High-accuracy calibration system laboratory table and enclosure.

washing the contaminated element was frequently required. The problem was greatly reduced by the addition of the air ionizer, placed in the air flow below the HEPA filter.

To prevent light transmission, the enclosure's side and top panels are made of black plastic. The inside of the panels is painted with a flat-black paint and textured to reduce specular reflections. Between the optical source and the detectors is a shield against scattered light. The shield is made of aluminum that was sanded and painted with a flat-black paint. The shield prevents the light scattered by the elements of the optical source from reaching the detectors. It also provides an approximately uniform dark field with constant temperature around the beam. The top of the enclosure over the LOC can be removed, so that the liquid-helium transfer line can be inserted. When not in use, the liquid-nitrogen transfer line (not shown) hangs from the top of the enclosure; nitrogen fills can be performed with the enclosure's top in place.

The enclosure panels behind and to one side of the LOC are left open, to facilitate routing the numerous cables onto the lab table, and to provide a passage for the clean air to escape without pressurizing the enclosure. The airflow out of the enclosure caused by the HEPA filter is not noticeable, but is strong enough to prevent dust from entering the opening. The HEPA filter runs continuously, except for when the minor draft it creates causes excess instability in test detectors.

Not shown in Figure 3 is a small ventilation hood, which was recently added to the enclosure panel above the LOC. The hood draws the gasses released by the boiling cryogenics out of the laboratory. Significant oxygen depletion can occur when the large amounts of nitrogen and helium gas are released during LOC cryogen filling, which can be a serious safety hazard for the operator. Therefore the oxygen content of the air is monitored during cryogen fill, to warn of any oxygen depletion that may occur. Air

flowing out of the ventilation hood is exchanged with fresh air from the building's ventilation system, so the gas concentrations return to normal atmospheric values after about an hour. A second independent ventilation system with flexible ducts was also recently added. Located over the optical source plate, it provides ventilation for air-cooled lasers.

Electrical power for the lab is isolated from the building by use of filtered surge suppressors, placed before and after an uninterruptible power supply (UPS). A "true" UPS is used, so that equipment with linear power supplies can be run from its inverter; the device continually runs its inverter, which produces an approximately sinusoidal output. Devices called "back" UPSs should be avoided, because some provide a simple square wave output when running from battery, which is sufficient for switching power supplies like those in a computer, but can damage devices with a linear power supply. Also, all back UPSs have a delay when switching from line power to battery power. The UPS used in the system also employs an internal isolation transformer. The isolation transformer helps filter power line noise by breaking the conductive path for low frequency noise and regenerates the line-neutral to earth-ground connection.

Most of the laboratory's electrical systems are powered by the UPS, except for devices that have a large power consumption, such as the computer monitors, HEPA filter, turbomolecular pump, hot air gun, and some of the lasers; these devices are protected only by surge suppressors. All of the electronics (except the turbomolecular pump, hot air gun, and some lasers) are connected to a single power circuit and ground to help reduce noise. It is not intended that the entire laboratory be run for long periods of time off the UPS's battery, but the UPS does keep the system from resetting during short outages. Also, the UPS can provide continuous power for devices that need it, such as the ion pump, during longer outages. Filters in the surge suppressors and UPS help keep high-frequency conducted noise out of the laboratory electronics.

Electrical shielding is provided to help reduce the amount of electrical noise in the system. The LOCR bridge electronics are isolated from the common electrical ground and its cabling is well shielded. Laboratory electronics used for test detector output measurement are also well shielded; for example, the system's digital multimeter (DMM) utilizes a guard shield that is switched by the three-pole multiplexer and propagated to analog voltage devices using shielded twisted-pair cabling. The multiplexer can also apply earth ground, or ground straps located on the laboratory table can be used when necessary. The effectiveness of the shielding is evident in the excellent measurement results produced by the system, despite the electrical noise present in the lab.

3.1.3 Mechanical systems: The lab table layout is shown with greater detail in Figure 4. The optical source resides on the elevated optical source breadboard, which is supported by five manually operated leveling posts. The posts provide accurate height adjustment, with a resolution of 0.25 mm. The LOCR is bolted to a small raised platform mounted on the detector breadboard. The raised platform reconciles the height difference between the two breadboards.

The detector breadboard is attached to a large, stepping-motor driven, ball-bearing linear translation stage. The stage has a load carrying capability of 90 kg and the load (not including test detectors) is approximately 75 kg. The LOCR itself weighs approximately 50 kg, when fully loaded with cryogenics.

Significant backlash in the detector stage's motion can occur. To minimize the potential backlash, motions are always concluded in a preferred direction. When motion in the opposite direction is required, the target location is over-shot by a few millimeters and the stage moves back to the target in the preferred direction. The preferred direction is also the direction in which the stage's drive nut contacts its fixed mount, so that there is no change in the absolute position even if the drive nut's clamp loosens. The

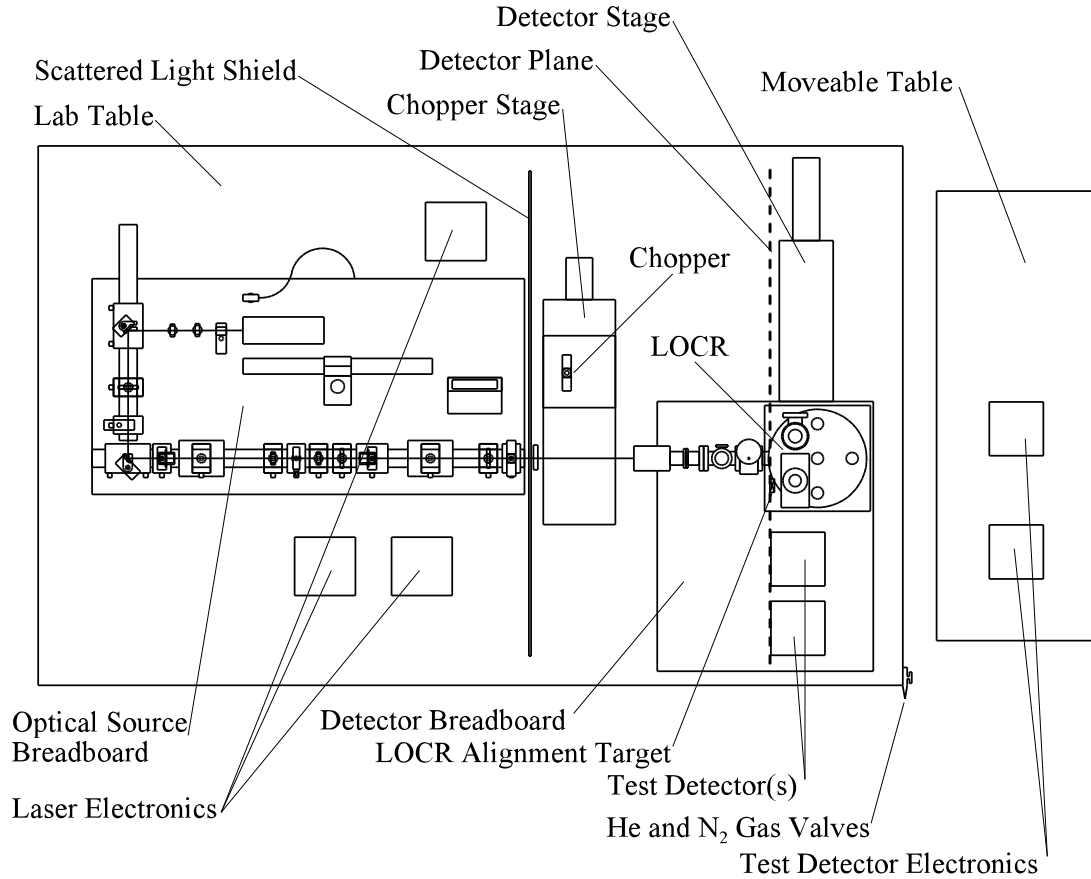


Figure 4. High-accuracy calibration system’s laboratory table layout.

resulting unidirectional reproducibility is approximately ± 0.04 mm, which is negligible compared to the alignment uncertainty, which is usually a few tenths of a millimeters or more.

A second motion-controlled stage, the chopper stage, is used to move an optical chopper into the beam’s path, for detectors that need chopped signals. A third and fourth stage, not shown in the figures, are combined in an x-y configuration. The x-y stage will be used to perform experiments such as scanning-slit beam diameter measurements, pinhole scans to measure the beam profile, precision detector centering, and detector spatial uniformity scans. The process control software does not currently support the x-y stage.

To smooth their motion, all of the motion-controlled stages used in the system are driven with a ramping acceleration. The maximum speed and acceleration time are set with user-accessible options in the process control software. The detector stage’s ramp time of 4 s is long enough to prevent excessive sloshing of the liquids in the cryogenic reservoirs, without resulting in motions so slow that the stage takes excessive time to translate.

Laser power supplies are placed on the lab table around the optical source breadboard. During calibrations, a movable table is placed behind the lab table, as shown in Figure 4. Detector electronics are placed on the movable table. The movable table must be moved out of the way during liquid-helium fills and turbomolecular pump operation.

3.1.4 Gas and cryogen handling systems: Both liquid and gaseous He and N₂ are used in the High-Accuracy Calibration System. The cryogenic liquids are used to cool the LOCR. The N₂ gas is used during LOCR pre-cooling and as a dry gas for venting and general cleaning purposes. The He gas is used primarily for pressurizing the liquid-helium container when liquid is transferred.

Valves with quick disconnect fittings for the He and N₂ gasses are conveniently located on the corner of the lab table, as shown in Figure 4. A variety of extension lines with quick disconnect adapters can be attached to the valves. Extensions include a high-pressure nozzle with trigger, a fitting designed to attach to the liquid-helium container for pressurization, and a line with knurled adapter for attachment to the rubber and plastic tubing used during LOCR pre-cooling and turbomolecular pump venting. Flexible plastic tubing is used in the gas delivery system; the tubing is purged with clean gas before use. The quick-disconnect fittings contain internal valves that close when the adapter are not connected, which helps prevent contaminants from entering the gas system. The fittings and valves can be operated while wearing insulated cryogenic gloves, which are worn whenever the cryogenic liquids are being handled.

The gas systems use high-pressure industrial-grade gas cylinders. Pressure regulators that are optimized for the delivery of gas at relatively low pressure (200 kPa maximum) are used with the cylinders. The gaseous He pressure regulator uses a metal diaphragm to prevent He leaks inside the regulator, and is capable of regulating a small negative pressure, so that the liquid-helium temperature can be lowered by applying a small vacuum to the LOCR's liquid-helium reservoir. The He gas will slowly permeate through the plastic tubing, so a vacuum is generated in the line when the valves at each end are closed. Valves that isolate the regulator from the plastic tubing are therefore always closed when the gas is not in use, to prevent the regulator from being damage by the vacuum. When not in use, both gas cylinder's main valves are always closed and the pressure in the regulators released, for safety and to protect the regulators.

The LOCR requires liquid helium and nitrogen, which are stored in commercial cryogenic liquid storage containers. The liquid-nitrogen container holds 160 l and the liquid-helium container holds 100 l. When performing calibrations, the liquid nitrogen is consumed in about two weeks and the liquid helium is consumed in about one month. The lead time for obtaining full containers is about two days for liquid nitrogen and two weeks for liquid helium. The liquid nitrogen is delivered through a flexible, insulated transfer line about 6 m long, which has a device called a vapor-phase separator at the LOCR's end. The phase separator allows the N₂ gas to escape while the liquid-nitrogen pours into the LOCR. The pressure of the boiling liquid inside the storage container is sufficient to push the liquid-nitrogen through the transfer line at an adequate rate. The liquid helium is delivered through a short, rigid, vacuum-insulated transfer line. The liquid-helium storage container is pressurized using He gas, to move the liquid more quickly through the transfer line. The liquid-helium transfer rate is closely monitored and controlled by the application of He gas pressure, because the transfer becomes increasingly inefficient if the liquid is transferred too slowly, or if the storage container is over-pressurized.

3.1.5 Optical system: The stabilized laser source is optimized to provide laser power that is very stable with time. The laser source resides on the optical source breadboard, shown in detail in Figure 5. The laser beam is vertically polarized and is delivered to the detector plane in an approximately Gaussian beam. The detector plane, shown in Figure 4, is where the detector's entrance or limiting aperture is placed.

Some beam quality is sacrificed in favor of power stability. The laser beams have high coherence, so to prevent interference from occurring, optical feedback must be avoided. The optical components are tilted to prevent back-reflections from reaching the laser and to prevent the formation of interferometer cavities within and between components. The tilt causes additional coma distortion in the lenses, but the

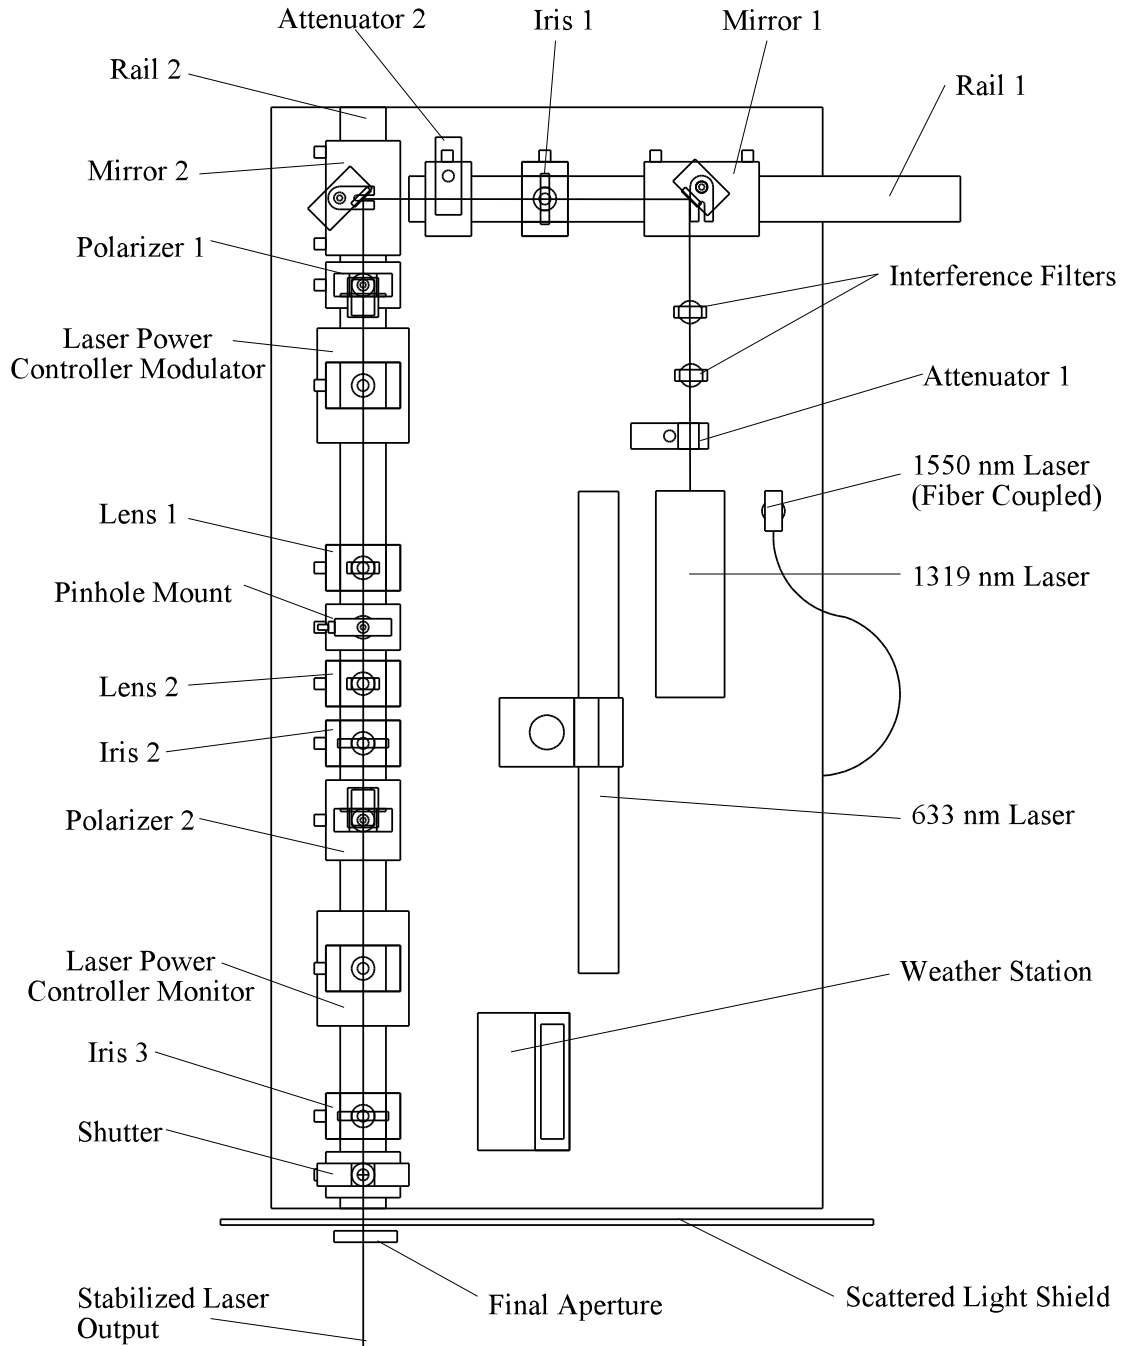


Figure 5. Detailed view of the optical source breadboard.

additional distortion is preferable to the power instability. Spherical doublet lenses are used to reduce the distortion caused by the tilt, and are carefully adjusted so that the beam passes through their centers, thereby minimizing the distortion. Other components, such as the polarizers and attenuators, are also tilted to prevent back-reflections. The polarizer tilt is performed carefully to prevent exceeding the polarizer's acceptance angle. Parallel-plate beam splitters are used with the power monitor elements, but the beam splitters are relatively thick, and are placed in the beam at an angle of about 45° , so that

interference does not occur in them. Thin optical elements are avoided, because even when tilted, a thin element can still produce interference fringes.

The lasers used in the stabilized source must meet several criteria. Ideally, the lasers should produce a linear polarization that is stable with time, a Gaussian laser beam, a single longitudinal mode or at least a single spectral line, and a stable or slowly drifting power. A stable, linear polarization is desired because polarization angle changes will become amplitude variations when the beam passes through Polarizer 1; lasers with other polarizations would be acceptable so long as the component with a linear polarization has a constant power. Lasers that produce a Gaussian beam are desired because Gaussian beams can be accurately modeled and are fully described by specifying the beam's waist location and diameter. Free-space lasers that produce only the TEM_{0,0} mode and laser beams delivered through a single-mode optical fiber produce a Gaussian beam of adequate quality. But if an optical fiber is used, the fiber must preserve the beam's polarization and a collimating lens is needed. Lasers with a single longitudinal mode or at least a single spectral line are desired, but filters can be used to clean up the spectrum of some sources with multiple spectral lines. Lasers that produce a stable or slowly varying power are desired because the laser power controller (LPC) cannot compensate for very large changes in beam power or for high-frequency fluctuations.

The lasers currently used produce approximately Gaussian beams with $1/e^2$ intensity diameter of about 0.25 to 1.0 mm at the laser's output aperture and a divergence angle of a few milliradians. Larger or more divergent beams can be significantly clipped by the circular entrance aperture, 4 mm in diameter, on the LPC modulator and monitor, resulting in undesired diffraction artifacts. As shown in Figure 5, the 1550 nm laser uses an optical fiber with an attached collimating lens. The others lasers shown have a free-space output. The layout shown in Figure 5 is typical, but other lasers are available, and when desired, a longer beam path can be constructed using additional mirrors.

The optical system's polarizers are calcite, so a very good extinction ratio of about 100,000 to 1 is obtained. Polarizer 1 is a Glan-Laser calcite polarizer, which can withstand a relatively high peak power, but has a narrow and asymmetric acceptance angle that is increasingly narrow at infrared wavelengths. To ensure that Polarizer 1 is placed so that the incident beam is close enough to normal incidence, its back reflection is traced all the way back to the laser's output aperture and the polarizer is angled just enough to prevent its back reflection from entering the laser. Polarizer 2 is a Glan-Thompson calcite polarizer, which has a larger acceptance angle, so it can be angled more than Polarizer 1 but it cannot withstand as much power as the Glan-Laser polarizer.

Most of the components are mounted on carriers along rails 1 and 2. The carriers provide smooth translation along the length of the rail, without rotation or translation in other directions. The carriers used with the lenses and irises have a built-in translation stage that provides fine alignment control in the horizontal plane, perpendicular to the length of the rail, without rotation or other translation. The optical components are mounted to the rail carriers using post holders, which allow the components to be rotated without translation. The two rails and the laser sources are at right angles to each other, and mirrors 1 and 2 are used with an incidence angle of 45° to help reduce mixing of the beam's polarization. Protected silver front-surface mirrors are usually used, but dielectric mirrors are also available when more efficient transmission in the visible spectrum is needed. The beam is directed with the optics so that it travels over the center of the rails at a constant height between components. Other components, such as the laser mounts, attenuator 1, and any interference filters are mounted directly to the optical source breadboard.

Attenuator 1 and any interference filters are placed in front of the laser, mounted directly to the optical source breadboard. To limit the amount of power at the filters, the majority of the attenuation is performed with attenuator 1, before the interference filters. Attenuator 2 is used to further reduce the

power to about 2 mW at the input to the LPC modulator. Attenuator 1 uses neutral-density reflective filters, and attenuator 2 uses absorbing glass filters; both are placed in the beam at a small angle to prevent interference.

Lenses 1 and 2 and the mounted pinhole form a spatial filter. The lenses are tilted slightly to prevent back reflections from entering the LPC modulator. The spherical, acromatic-doublet lenses are positioned so that the beam passes through the approximate center of each lens, to reduce the coma distortion caused by the tilt; the lenses can be used slightly off-center to steer the beam as necessary to compensate for imperfections in Polarizer 1 and the LPC modulator. The two lenses are tilted in opposite directions to prevent a horizontal translation of the transmitted beam. Acromatic-doublet lenses are used instead of single-element lenses so that there are more refracting surfaces, to reduce spherical aberration over a wider range of wavelengths. Lens 1 is placed at a distance from the LPC modulator to minimize the tilt angle required to prevent the back-reflection from re-entering the modulator, but a tilt of a few degrees is usually necessary because the spherical surface of the lens causes the back reflection to diverge. The spherical lenses require less tilt and produce less aberration with the same tilt than comparable single-element parabolic or “best form” lenses.

Several lens pairs, each with a diameter of 2.54 cm, are available with focal lengths of 50, 75, and 100 mm. The lenses are used throughout the visible to near infrared spectrum. The lenses have a single layer of antireflection coating optimized for visible wavelengths, but they have sufficient transmission to allow operation at infrared wavelengths. Irises 1 to 3 are used during beam alignment and are used to block back reflections and scattered light after alignment, but are **not** used to restrict the extent of the Gaussian beam (which would produce diffraction artifacts).

In the high-accuracy calibration system, the spatial filter [7], which consists of Lens 1, Lens 2, and the pinhole, is used somewhat differently than usual. The beam illuminating Lens 1 is an approximately Gaussian beam, produced by the laser source. Therefore, Lens 1 causes a Gaussian beam waist to form at a location slightly behind the Fourier plane. A pinhole aperture is placed at the beam waist to effectively remove any higher spatial frequencies present in the incident beam. Scattered optical power, which surrounds the approximately Gaussian beam, is effectively removed by the spatial filter. The spatial filter does not create a Gaussian beam from a non-Gaussian source, it helps remove only non-Gaussian components.

The pinhole used in the spatial filter has a diameter at least twice the size of the $1/e^2$ intensity diameter of the Gaussian beam waist, so that the intensity of the light at the edge of the pinhole is small and the resulting diffraction produces minimal artifacts in the reconstructed beam. In practice, the result is that the scattered power around the incident Gaussian beam is effectively removed, but the Gaussian beam itself is essentially unchanged. This modified spatial filter can significantly improve the beam quality, since there is usually significant scattered power produced by the mirrors, polarizer, attenuators, and the laser source itself. But the quality of the Gaussian beam emanating from the modified spatial filter is therefore only as good as the quality of the Gaussian beam produced by the laser source itself, so it is important that the laser source itself is aligned well. The optical system can degrade the Gaussian beam's quality, but does not improve it, other than by removing the non-Gaussian components or scattered light.

To control the size and divergence of the beam in the detector plane, the location of Lens 2 is also varied from that in the classic spatial filter. The lens focal lengths and the distance from the laser source to Lens 1 can also be varied to obtain other beam sizes and divergence angles, but varying the location of Lens 2 provides a fine control for the size and divergence of the resulting beam. Because the beam produced is a good approximation to the ideal Gaussian beam, the behavior of the beam is accurately predicted using

the Gaussian beam model described in Appendix D. The size of the beam waist formed between the lenses is also determined using the model.

The optical shutter uses a plastic-coated blade that opens and closes quickly, with rise times and delays that are less than 10 ms. The amount of time the shutter is open is measured with the exposure counter, which is a precision time-interval counter that is accurate to better than 0.001 % (for time intervals in seconds). Asymmetry in the shutter switching time is not accounted for in the measurement, but is usually negligible compared to the period during which the shutter is open. The period in which the shutter is open is important for energy meters only, where the energy is the average power multiplied by the period the shutter is open.

The final element in the optical source is a circular final aperture. The final aperture is mounted to the shutter's housing, so that its location relative to the laser beam is constant when the height of the optical source is varied. The final aperture is a 4 mm diameter, painted ceramic aperture; it is used primarily to block a stray beam that is produced by Polarizer 2. The final aperture must be accurately aligned, or it can significantly clip the laser beam and cause diffraction artifacts in the detector plane. A larger final aperture can be used to accommodate other beam diameters, but an aperture much larger than 4 mm may not block the stray beam. The stray beam diverges from the main beam at a small angle, so the final aperture must be small and placed sufficiently far away from Polarizer 2 so that the stray beam diverges enough to be blocked by the aperture, hence its mount extends to the detector's side of the scattered light shield. The stray beam is visible when the 633 nm laser is used; it may not be detectable with IR cards at infrared wavelengths, but probably is still present. If the final aperture were not present, the stray beam's divergence is typically such that it would be captured by a detector with a limiting aperture of 10 mm diameter, but missed by a detector with a limiting aperture of 5 mm diameter. This effect would be compensated for by the relative aperture transmittance correction, but the measurement uncertainty is lower when the problem is eliminated by using the final aperture. Additionally, the stray beam may have a different polarization from that of the main beam, which would complicate the LOCR Brewster's angle window transmittance and the measurement uncertainty analysis, if the beam were not blocked.

Because it is not much larger than the beam's diameter, the final aperture does produce diffraction artifacts in the detector plane, but the minor diffraction artifacts are preferable to the stray beam that would otherwise be present. To minimize the diffraction artifacts, the final aperture's diameter should be at least twice as big as the $1/e^2$ intensity diameter of the beam at the aperture. Conversely, the small final aperture can help block diffraction artifacts and scattering produced by the elements in the system that are after the modified spatial filter. The precise effect of the final aperture on the beam profile in the detector plane can be studied in detail using precision beam-profile scanning techniques, such as that recently developed for measuring the mode field diameter of an optical fiber [8].

3.1.6 Laser power stabilization system: Laser power stabilization is provided by the laser power controller (LPC). The LPC consists of three separate parts; the LPC modulator, the LPC monitor, and the LPC electronics. The LPC modulator uses a liquid-crystal device to rotate the polarization of the transmitted laser beam, by up to 90°. Polarizer 2 demodulates the beam, converting the polarization modulation to an amplitude modulation. Polarizer 2 also removes any polarization mixing caused by the intervening lenses of the spatial filter, ensuring that a strongly polarized beam is delivered to the LOCR Brewster's angle window. The LPC monitor is placed after Polarizer 2, so that it measures the power in the demodulated beam. A parallel-plate beam splitter in the LPC monitor samples a small portion of the beam and the sampled beam's power is measured using an internal detector. The monitor's power measurement is fed to the LPC electronics, where a control system produces the modulator signal. The LPC monitor's beam splitter and detector combination are calibrated in the factory, so that the power measured by the detector can be related to the absolute power in the transmitted beam.

The LPC modulator and LPC monitor are each housed in identical, precision-engineered enclosures. The beam passes through the devices with no appreciable horizontal offset or angular deflection, but a small vertical offset is produced. The enclosures have fixed circular apertures: a 4 mm diameter entrance and 5.5 mm diameter exit apertures. The beam must be small enough at these apertures to keep any resulting diffraction artifacts to a minimum, so the beam should have a $1/e^2$ intensity diameter of one-half the aperture diameter or less. This requirement is usually easily met at the LPC modulator, where the Gaussian beam has a relatively small diameter; but the size of the apertures on the LPC monitor, which is much closer to the detector plane, can limit the diameter of the largest usable beam. The LPC modulator and monitor can however produce a slight wavelength-dependent angular deviation in the vertical plane. The offset produced by the LPC modulator is compensated for by slight steering with Lens 1. The deviations are small, typically a few milliradians.

Two complete LPC systems are available in the high-accuracy calibration system. The visible wavelength LPC uses Si detectors to measure the power and is designed to operate at wavelengths from 425 to 780 nm, with input laser power from 0.5 mW to 4 W. The infrared wavelength LPC uses InGaAs detectors and is designed to operate at wavelengths from 950 to 1,700 nm, with input laser power from 0.5 mW to 2.0 W. Each modulator has a maximum optical transmission of 80 to 85 %, within the stated wavelength range. The LPCs can be used outside their specified wavelength range, but the optical power reported by the LPC electronics will be inaccurate and the stabilizer's performance will be reduced. Therefore, operation over the wavelength range of 780 to 950 nm is possible with the available equipment, albeit at reduced performance.

The LPC monitor uses a parallel-plate beam splitter, set at an angle of 45°. The beam splitter is oriented so that it is near Brewster's angle for the transmitted beam's polarization, so very little of the power is reflected to the monitor detector. The beam splitter is thick enough (a few millimeters) to prevent interference from occurring. The small sample ratio limits the LPC to a usable minimum power of about 100 μ W, lower powers become increasingly noisy because of noise in the monitor detector. Powers lower than 100 μ W can be achieved without adding excessive noise, by mounting the power monitor so that it is rotated 90° around the beam's propagation axis, then a greater percentage of the transmitted beam is reflected to the monitor detector, improving its signal-to-noise ratio. The upper power range of the LPC is ultimately limited by the damage level of the LPC modulator's liquid crystal, but in practice the lasers used limit the power to about 10 mW. Both the LPC modulator and monitor have internal neutral density filters, which are used to attenuate the power reflected to the monitor detectors inside the modules, when high powers are used. The LOCR cannot be used with powers much greater than 1 mW, but transfer standards could be used with higher powers.

The spectral response of the monitor detectors used in the LPC is calibrated at the factory. The vacuum wavelength of the beam being stabilized is entered into the LPC electronics, where the monitor's spectral response is corrected. The resulting power measurement, which is accurate to within a few percent, can be displayed on the LPC electronics and read through its RS-232 port.

The LPC modulator can also function independently of the external LPC monitor. Each LPC modulator module contains an internal polarizing beam splitter and internal monitor detector, so can be used without the external polarizer and LPC monitor. The LPC modulator's internal monitor is used during source alignment, but to compensate for loss and polarization mixing in the spatial filter and loss in Polarizer 2, the external LPC monitor is used during calibrations. Polarizer 2 has a greater extinction ratio than the polarizer in the LPC modulator and is the last polarizer in the optical system, so it sets the polarization angle of the beam at the detectors. Polarizer 2 can therefore be used to adjust the polarization angle to match the orientation of the LOCR's Brewster's angle window. Using the current mount, rotating the LOCR's window mount is difficult and impossible to do with precision, while adjusting the beam

polarization with Polarizer 2 is simple. However, to minimize the power loss and maximize the dynamic range of the system, the polarization angle of the laser source, Polarizers 1 and 2, and the LPC modulator's internal polarizer should all be set as close as possible to the same angle.

The LPC electronics is a control system that is effective at compensating for drift in the input laser power, within its bandwidth of 5 kHz. Large noise spikes or drift can exceed the device's modulation range, so the LPC electronics package includes an error light that flashes when the modulator is driven out of range. When operating properly, the power drift is typically less than 0.05 % peak-to-peak per hour, for powers around 1 mW. The error light will illuminate continuously if the modulator is driven completely out of range and no regulation is provided. A computer can be used to check the LPC error status and transmission ratio, through the LPC electronic's RS-232 port, to ensure it is operating within its range.

3.1.7 The relative aperture transmittance correction: The high-accuracy calibration system's optical source is designed to deliver to the detectors a good approximation of the theoretical Gaussian beam. Probably the most significant difference between the actual and theoretical Gaussian beams is that the actual beam contains more power away from the beam's center than the theoretical Gaussian beam does. The spatial filter helps reduce the scattered power from around the Gaussian beam, but some remains and the subsequent optical elements can produce more. To compensate for this imperfection, the relative aperture transmittance correction, called T_{RA} , is applied.

A Gaussian beam has the majority of its power contained near the center of the beam and the beam's intensity drops rapidly at increasing distance from its center. Technically the beam has tails that extend to infinity, but they become very small after a few beam radii. If a detector's limiting aperture diameter is several times the size of the Gaussian beam's diameter, it is common to assume that the detector captures the entire beam. However in the real optical system, there is significantly more power in the tails than in the theoretical Gaussian beam, so the amount of power that is not captured by such detectors can be significant in a high-accuracy calibration. Detectors with entrance apertures that are different sizes or shapes can therefore be exposed to different absolute powers, even if their apertures are many times larger than the beam.

The correction factor T_{RA} compensates for the amount of power that is missed by the standard and test detectors, so that the resulting detector responsivity calibration factor is based on the amount of power that was actually delivered through each detector's entrance aperture. Usually, the customer specifies a beam diameter that is several times smaller than their detector's aperture, with the assumption that essentially all of a theoretical Gaussian beam will be captured. The correction T_{RA} is consistent with this assumption, and in this case, $|T_{RA} - 1|$ is a measure of the imperfection in calibration system's beam. But more generally, the correction results in the same calibration factor being obtained for given test and standard detector, independent of the beam's actual intensity profile, so long as both the standard and test detector's responsivity is uniform and linear within their apertures. In practice, the T_{RA} correction is usually negligible compared to the overall uncertainty in the calibration, unless the test and standard detectors have very different aperture sizes.

There is historical precedent for applying a correction for a detector's aperture transmittance. The relative aperture transmittance correction is similar in concept to the spillover monitor-based correction used with an old device, the NIST BB calorimeter [9]. In the BB's application, the total energy in the beam was desired. It was evident that the BB's optical receiver could miss a significant portion of the large multi-mode beam usually used with the calorimeter, so a spillover monitor (which has an entrance aperture much larger than the calorimeter's) was added to measure the amount of energy spilling over the calorimeter's aperture. Originally, the spillover monitor was used only as a quality control measure, to determine whether the calorimeter is capturing enough of the beam to meet its specified uncertainty. An

uncertainty was assessed for the expected spillover and the monitor was used to determine whether the actual amount of spillover was within the assessed uncertainty limit. The spillover was later corrected for, to obtain lower calibration uncertainties.

The aperture referred to in this discussion is ideally the detector's limiting aperture. But some detectors have no physical limiting aperture and the size of their entrance aperture can be deceiving. With such detectors, the optically responsive element inside can be smaller than the entrance aperture, so some light entering their entrance aperture misses the detector, even if the light rays are perfectly collimated. For these detectors, an equivalent limiting aperture called the apparent aperture is used. The apparent aperture of such a detector is shown in Figure 6; the detector's apparent aperture is clearly much smaller than its approximately 12 mm diameter circular entrance aperture. The apparent aperture is simply the 50 % contour line in a spatial uniformity scan of the detector's responsivity [10].

The 50 % responsivity contour is used because it inscribes the edge of the detector's physical limiting aperture, when one is present, so long as the beam used in the uniformity scan is small compared to the aperture's feature size. For example, when a scanning beam is centered over the edge of a relatively large limiting aperture, half will fall on the detector and half will be blocked, so the responsivity drops 50 % from when the beam is centered in the aperture. Since the real scanning beam is not infinitely small compared to the aperture, the measured apparent aperture will be somewhat smaller than the actual limiting aperture. When the Gaussian scanning beam's $1/e^2$ diameter is one-fifth the minimum aperture diameter or less, the reduction in the size of the measured apparent aperture is considered sufficiently small for the purpose of estimating a detector's T_{RA} , given that the uncertainty in the measurement of T_{RA} itself is relatively large. The process used to measure T_{RA} is described in Section 4.1.5 and Appendix E.

3.2 The LOCR Primary Standard

The primary standard used in the high-accuracy calibration system is a laser-optimized cryogenic radiometer, abbreviated as LOCR. Cryogenic radiometers provide the most accurate means of measuring absolute optical power that is currently available. The uncertainty of the primary standard's power measurement directly limits the uncertainty obtainable with the calibration system, so a cryogenic radiometer primary standard is necessary to achieve the desired high-accuracy detector responsivity calibration.

3.2.1 Basic cryogenic radiometer operating principal: Cryogenic radiometers gain their performance advantage over room-temperature atmospheric-pressure devices by improved diffusivity in the optical receiver, reduced radiation from the optical receiver, and reduced power dissipation in the electrical heater leads. These factors reduce the inequivalence between electrical and optical heating, thereby improving the accuracy of the electrical substitution. To achieve these properties with the materials commonly available today, the devices operate usually near the liquid-helium temperature of 4 K. Operation at cryogenic temperature necessitates the use of ultra-

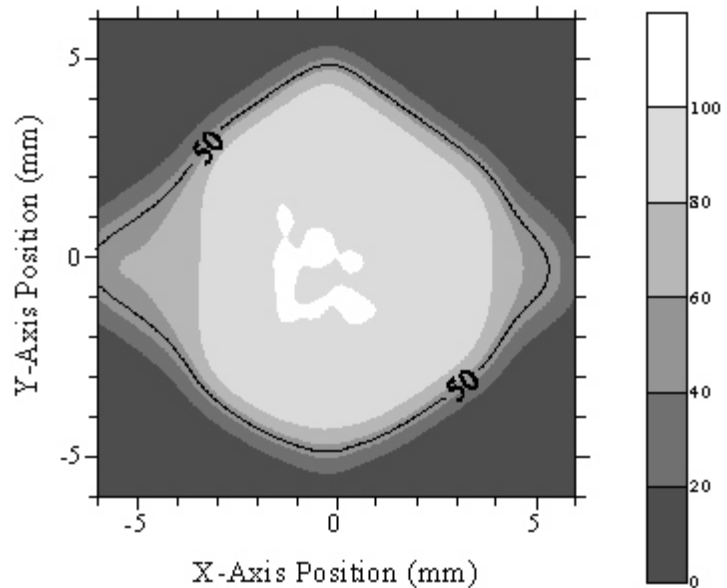


Figure 6. The apparent aperture of a detector that has no limiting aperture.

high vacuum inside the radiometer to insulate the cold interior from the external environment, which also eliminates any convective heat flow inside the device, but necessitates the use of a window that allows the light to enter the device.

Laser-optimized cryogenic radiometers such as the LOCR use an optical receiver design that optimizes the capture of the collimated laser beam and use a Brewster's angle window that reflects very little of the incident laser power. Polarized laser light passes through the Brewster's angle window with minimal attenuation and is absorbed in the optical receiver, where the optical power is converted to heat. The amount of heating power is then measured by electrical substitution, so the device is a primary standard, traceable to SI units through the electrical substitution.

3.2.2 Description of the LOCR: The LOCR has been previously described [6, 11, 12], so its internal construction will be discussed only briefly here. The LOCR is a custom designed system that was built by a commercial manufacturer [13]. An external view of the LOCR is shown in Figure 7, and Figure 8 shows an internal cross-section view. The LOCR uses both liquid-helium and liquid-nitrogen, held in reservoirs that are exposed to atmospheric pressure. The hold time for liquid nitrogen is approximately 14 hours, so the nitrogen is refilled at least twice a day. The hold time for liquid helium is approximately 2.5 days, so the helium reservoir is filled twice a week (the cryogenics expire over the weekend). The cryogenic liquid hold times are affected by room temperature, and if the room is very warm, the hold times are reduced. Vacuum inside the cryostat is provided by the cryo-pumping produced by the cold temperatures, which is enhanced by the addition of the carbon getter attached to the liquid-helium temperature heat shield, and by an ion pump that is located on top of the cryostat. A partial vacuum can be applied to the helium reservoir to further reduce the temperature of the liquid, but the device has not yet been operated in this mode, and the liquid-helium's hold time would likely drop considerably.




Figure 7. External view of the LOCR.

As per the manufacturer's recommendation, the LOCR is brought back to room temperature weekly. Leakages through the o-ring seals in the cryostat, valves, and window mount have the potential for creating ice on the optical receiver, thus changing its optical properties, so the LOCR is warmed to room temperature weekly to prevent any significant ice buildup. The warming also clears the carbon getter and ion pump, allowing lower pressures to be maintained. Allowing the LOCR to warm up over the weekend also eliminates the need for an automatic liquid-nitrogen filling system, or a technician coming in twice a day to refill the nitrogen manually.

The dashed line in Figure 8 shows the location of the detector plane inside of the LOCR. The LOCR has three circular apertures in front of its optical receiver. The detector plane is located at the final aperture, which is the smallest (10 mm in diameter) and the closest to the optical receiver. The optical receiver itself is slightly larger, so the final aperture is ideally the limiting aperture. By convention, the detectors under test are placed on the detector breadboard so that their entrance or limiting aperture is in the detector plane.

Small imperfections in the LOCR's construction and the optical system's alignment necessitate tilting the LOCR a few milliradians so that the beam passes through both the center of its window mount and the center of its optical receiver. The tilt results in the LOCR's apparent aperture being slightly smaller than its limiting aperture and elliptical; in its current configuration, the LOCR's apparent aperture measures 9.5 mm in diameter horizontally and 8.5 mm in diameter vertically.




Figure 8. Internal view of the LOCR.

Thermal modeling [11, 12] showed that when operating at the nominal receiver temperature of 5.4 K, the inequivalence between electrical and optical heating is negligible (0.0005 %), even when the optical heating power is applied at an extreme location on the receiver.

Figure 9 shows the LOCR's optical receiver, with the normal and extreme incident beam locations that were used in the model. Since it is unlikely that the optical power would be applied at such an extreme location, the actual inequivalence encountered is likely to be much smaller than that determined by the modeling. Superconducting lead wires are used with the LOCR's resistive heaters to eliminate loss in the heater leads, thereby improving the accuracy of the electrical heating power measurement.




Figure 9. The LOCR's optical receiver.

To aid in describing the electrical operation of the LOCR, a simplified schematic, shown in Figure 10, is used. As shown in the figure, the LOCR uses two precision digital temperature controllers, one to regulate the heat sink near liquid-helium temperature, and one to regulate the optical receiver at slightly higher temperature. The temperatures at the receiver and heat sink are sensed with germanium resistance thermometers (GRTs). The GRT resistance is measured using precision alternating-current bridge circuitry, located in the LOCR's base. The temperature-control electronics apply power through resistive heating in the thin-film resistors located on the heat sink and optical receiver, and use the GRTs for feedback. The heat-sink temperature is regulated because the temperature of the liquid-helium well itself can change with liquid level and atmospheric pressure.

Application of the optical power is controlled using an external shutter. Electrical substitution in the LOCR is provided by measuring the drop in electrical heating power that is required to maintain the optical receiver's temperature, when the optical power is applied. This mode of operation is usually called the active or dynamic mode. The optical power is therefore the difference between the two absolute electrical powers; the electrical power when no optical power is applied (the background power) and the electrical power when the optical power is applied (the illuminated power). The uncertainty in the two electrical power measurements themselves can be larger than the uncertainty in the difference, because some of the uncertainties in the power measurements are correlated, as described in Appendix F. In practice, the background and illuminated electrical powers are derived using the dual-baseline measurement, described in section 4.3.1, but the underlying concept is the same.

The active mode of operation is advantageous because the optical receiver is maintained at the same temperature, so its temperature and therefore its thermal properties do not change when the optical power is applied. And because the

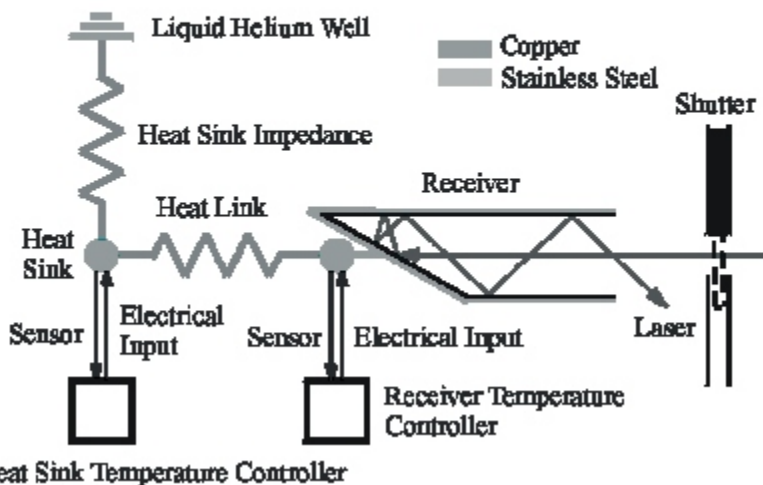


Figure 10. General functional schematic of the LOCR.

receiver's temperature ideally never changes, the measurement of optical power can be performed relatively quickly. The temperature controllers also have an “open” mode of operation, where they simply apply a fixed amount of heating power. Open mode is used during LOCR electrical calibration, but the LOCR is not intended to operate in open mode during optical power measurement.

One disadvantage of the active operating mode is that the device's overall performance is directly tied to the performance of its temperature-control systems, so temperature-control systems of very high quality are required. The LOCR uses digital control systems that can track the step change that occurs when the optical power is applied with no offset. At its typical operating temperature of 5.4° K, the thermal time constant of the optical receiver itself is approximately 10 s.

Since the LOCR is a thermal device that measures essentially any radiation, changes in the background temperature can affect the desired optical power measurement. Experiments have shown that the LOCR's temperature coefficient is approximately $-950 \text{ nW/}^\circ\text{C}$ [11, 12], due to radiative heat transfer from the window assembly to the optical receiver. Specifically, this is the change in receiver heater power that is necessary to maintain the receiver at a constant temperature, for a given change in the window assembly temperature. Since the window assembly is not temperature regulated, it follows room temperature; so the above is essentially the room-temperature coefficient. Fortunately, a linear drift in the window assembly temperature is compensated for by using the dual baseline and the power interpolation methods used with the calibration system.

Other environmental influences also affect the LOCR, but are usually negligible compared to other sources of uncertainty in the system. For example, since the LOCR's cryogenic reservoirs are exposed to the atmosphere, a change in atmospheric pressure will cause a change in the temperature of the cryogenic liquid, which changes the temperature of the thermal shields connected to the reservoirs. The heat sink's temperature regulator isolates the optical receiver's conductive heat flow path from the change in reservoir temperature, but a thermal shield is connected directly to the liquid reservoir, so its temperature can change, thereby changing the radiative heat transfer from the shield to the receiver. Similarly, the temperature of the LOCR's vapor-cooled liquid-nitrogen temperature shield changes with the level of the liquid in the reservoir. And atmospheric humidity may change the window transmittance, as water vapor is adsorbed on the surface of the window. Environmental parameters during the calibration are logged by the weather station continuously, so that if other significant environmental effects are discovered later, a post-processing correction can be applied.

3.2.3 LOCR electrical calibration: The LOCR's temperature control electronics are calibrated in terms of the amount of electrical power that they apply to the LOCR's thin-film heaters. Imperfections in the electronics cause an error in the electrical power that is reported by the temperature controller's digital electronics. To correct for the error in the reported power, the actual power delivered is measured independently and a correction factor for the digital power reading is stored in the temperature control electronics.

A simplified schematic of the LOCR's electrical heating circuit is shown in Figure 11. The resistor R_s is the standard resistor, which has a known resistance. The resistor R_H is the thin-film resistor placed on the LOCR's optical receiver, through which the electrical heating power is delivered to the receiver. The voltage V is applied by the temperature control electronics and the current I flows through the resistors. It is unimportant whether the temperature controller actually uses voltage or current sources, since either will result in the voltage V across the resistors when the current I flows through them. During electrical calibration, the




Figure 11. Simplified LOCR electrical calibration circuit.

voltage drops across the resistors V_S and V_H are measured using the calibration system's precision multimeter. The power dissipated in the resistor R_H , called P_H , is given by:

$$P_H = V_H \frac{V_S}{R_S}. \quad (1)$$

The power P_H determined using the measured voltages is then compared to the value of P_H returned by the temperature controller's digital electronics; the difference is the correction factor for the digital electronics. During calibration of the LOCR's electronics, the correction factor for a large array of applied powers is determined. The correction factors are analyzed and stored in the LOCR's electronics and subsequent measurement of electrical power are automatically corrected using the stored data. To update the stored correction factors, the entire array of powers must be measured again, which is very time consuming.

The stored calibration factors are accurate when the electronics are used in an environment similar to that present during the calibration and for some period of time after the calibration. The non-zero temperature coefficient of the electrical components cause the digital power reading to vary with temperature, and humidity can affect some of the electrical components, causing a sensitivity to the environment. And some of the electrical components, such as the precision resistors, can change with time and usage as they age. Therefore, for the electrical calibration to be accurate, the calibration must be performed in an environment similar to that which will be present during optical power measurements, and the measurements must be performed a relatively short time after the electrical calibration (a 90-day limit on the age of the electrical calibration is assumed in the uncertainty calculations).

Since the LOCR electrical calibration process is very time consuming, it is not efficient to perform a full electrical calibration for each optical calibration that is performed. Instead, a relatively quick calibration check is performed during each test detector calibration, at only the power levels of interest in the test detector calibration.

In the calibration check, the power P_H is measured as above, but the measured power is compared to the corrected digital power, using the correction factor stored in the electronics from the previous full electrical calibration. The difference between the measured and corrected digital powers is then the error in the corrected digital power reading. Instead of using the resulting error to update the correction factor stored inside the digital electronics, the error is simply recorded. Using this calibration check method, the error in the digital power readings is determined at only the power levels that will be used in the forthcoming high-accuracy calibration and the resulting error is corrected in post-processing.

The calibration check must be performed at each nominal electrical power that is applied to the LOCR's receiver during the optical calibrations. For example, the nominal receiver power during background illumination (no laser power) is 2.0 mW, so a calibration check at 2.0 mW is needed. A calibration check is also needed for each laser power that will be used in the detector calibration. For example, if the test detector is to be calibrated with laser powers of 1.0, 0.5, and 0.1 mW, the calibration check must also be performed for receiver powers of 1.0, 1.5, and 1.9 mW, respectively.

Because of the LOCR's room-temperature sensitivity, the receiver temperature that results from the nominal 2.0 mW of electrical power can change somewhat with room temperature. Because the electrical power is calibrated at only the nominal power levels, the nominal electrical power that is delivered to the receiver is given priority over the specific receiver operating temperature. The receiver regulation temperature used by the temperature control is therefore determined by applying the nominal

receiver electrical power using the temperature controller's open mode, allowing the receiver to reach its steady-state temperature, and measuring the resulting temperature. Then, the temperature controller is set to maintain the measured temperature. Fortunately, the room temperature in the laboratory does not vary much from the nominal room temperature of 23° C, so to date the resulting receiver temperature has been in the range of 5.405 ± 0.033 K, which is considered acceptable since the thermal properties of the receiver, modeled at 5.4 K, will not change significantly over this range. The nominal 2.0 mW of electrical power that is applied to the heat sink can be adjusted to optimize the receiver temperature that results from application of the nominal receiver heater power, but such an adjustment has not been needed to date. If unusual room temperatures are present, such as if the room temperature is intentionally varied to study the test detector's temperature coefficient, adjusting the power applied to the heat sink may be necessary to keep the receiver temperature within a reasonable range.

The heat sink and optical receiver temperature controllers are essentially identical, so both are calibrated the same way. However, only the optical receiver's temperature control must be accurately calibrated, because the absolute power applied to the heat sink is irrelevant to the measurement of optical power. Therefore, it is not necessary to check the heat sink temperature controller's calibration during test detector calibration. The heat-sink temperature controller's calibration is however checked empirically when the LOCR is operated; a large change in the heat sink's temperature for a given applied power would indicate a problem.

The LOCR's design was optimized for optical power measurements of about 1 mW, but the device can be used over a range of powers, with a corresponding increase in its uncertainty. Optical receiver power levels much above 4 mW cannot be used because the resulting receiver temperature will exceed the superconducting wire's critical temperature. The conductance of the heat link which connects the optical receiver to the heat sink controls the magnitude of the receiver's temperature rise over the heat sink temperature, for a given receiver power. Unlike some other cryogenic radiometers, the LOCR's heat links are not intended to be changed, so cannot be easily tailored for operation at other power levels.

3.2.4 The NIST Brewster's angle window: The LOCR Brewster's angle window, shown in Figure 12, was designed and built by NIST. The Brewster's angle window is mounted independently from the LOCR, on the window assembly, as shown in Figure 7. The window assembly rides on a transverse rail, which is attached to the detector breadboard, in front of the LOCR. The rail allows the window assembly to be precisely moved in front of the LOCR, with little tilt or other undesired translations. The rail is used to move the window assembly away from in front of the LOCR, for window transmittance measurement.




Figure 12. The NIST Brewster's angle window mount.

The window mount is connected to a kinematic gimbal mount, that allows precise angular adjustment of the window along vertical and horizontal axes. The window mount gimbal also holds the NW 40 clamping vacuum flange. The clamping flange allows easy removal of the window from the LOCR. Together, the window mount, window mount gimbal, rail carrier, and vacuum flange are called the window assembly. The window assembly is rigid and heavily articulated, with adjustments for setting the mount's height as well as the angle of the Brewster's angle window and vacuum flange location.

The window mount was designed to minimize the mechanical stress placed on the window, thereby reducing any stress-induced birefringence that may occur inside the window. It accepts a common 51 mm diameter, 10 mm thick circular window, with loose tolerances. The window's vacuum seal is provided by a thin, soft (50 durometer) nitrile o-ring, which was chosen because it compresses sufficiently under the approximately 18 kg of force exerted by atmospheric pressure when the back of the window is exposed to vacuum. No mechanical clamp is used; atmospheric pressure alone compresses the seal. Despite being adequately compressed, the o-ring has to be lightly lubricated with vacuum grease to reduce the leakage to acceptable levels. A shoulder machined into the mount, underneath the window, is intended to help limit any spread of the lubricant across the back side of the window by breaking the capillary action. A butyl rubber o-ring is preferable to the nitrile o-ring because butyl would provide a better seal with the relatively low compression, but the nitrile o-ring works adequately when lubricated. To prevent problems with multiple reflections, the inside of the mount is painted with a flat black, vacuum-compatible paint. The mount has a circular clear aperture with a diameter of 19 mm.

The window currently used with the LOCR is well polished but uncoated and made of quartz glass. The window is a circular plane-parallel window, 51 mm in diameter, 9.5 mm thick, with a wedge angle that is less than or equal to 5 min of arc (1.45 mrad), a surface flatness of $\lambda/10$ at 633 nm, and a surface quality of 10-5. The ability to rotate the window when it is on the mount is useful, because only a small section of the window is actually used. So if there is an imperfection in the glass in the area that is being used, the window can be rotated so that a better portion is used instead. However, the window should not be rotated while the mount is evacuated, because a catastrophic rush-in of air will result if the vacuum seal is broken. The mount will accommodate variations in the window's diameter of about 1 mm and any window thickness up to 10 mm.

Windows with smaller wedge angles are available at additional cost, but are unnecessary because the additional attenuation caused by the small wedge is compensated for in the window transmittance measurement. Surface quality and internal scatter centers typically have a much bigger impact on the transmittance than the small wedge angle. The wedge reduces the window's transmittance because the laser's incidence angle cannot be precisely Brewster's angle on both sides of the window, but the reduction in transmittance is accounted for in the window transmittance measurement. The change in the window's transmittance that occurs when vacuum is applied to the back of the window is accounted for in the window transmittance uncertainty.

The window mount also has a tight-fitting dust cover, shown in Figure 12. The cover helps protect the window from contamination when it is not in use and is used to aid in the window alignment, as described in Section 4.2.4. The dust cover limits the maximum thickness of the window to about 1 cm, because a thicker window will block the cover.

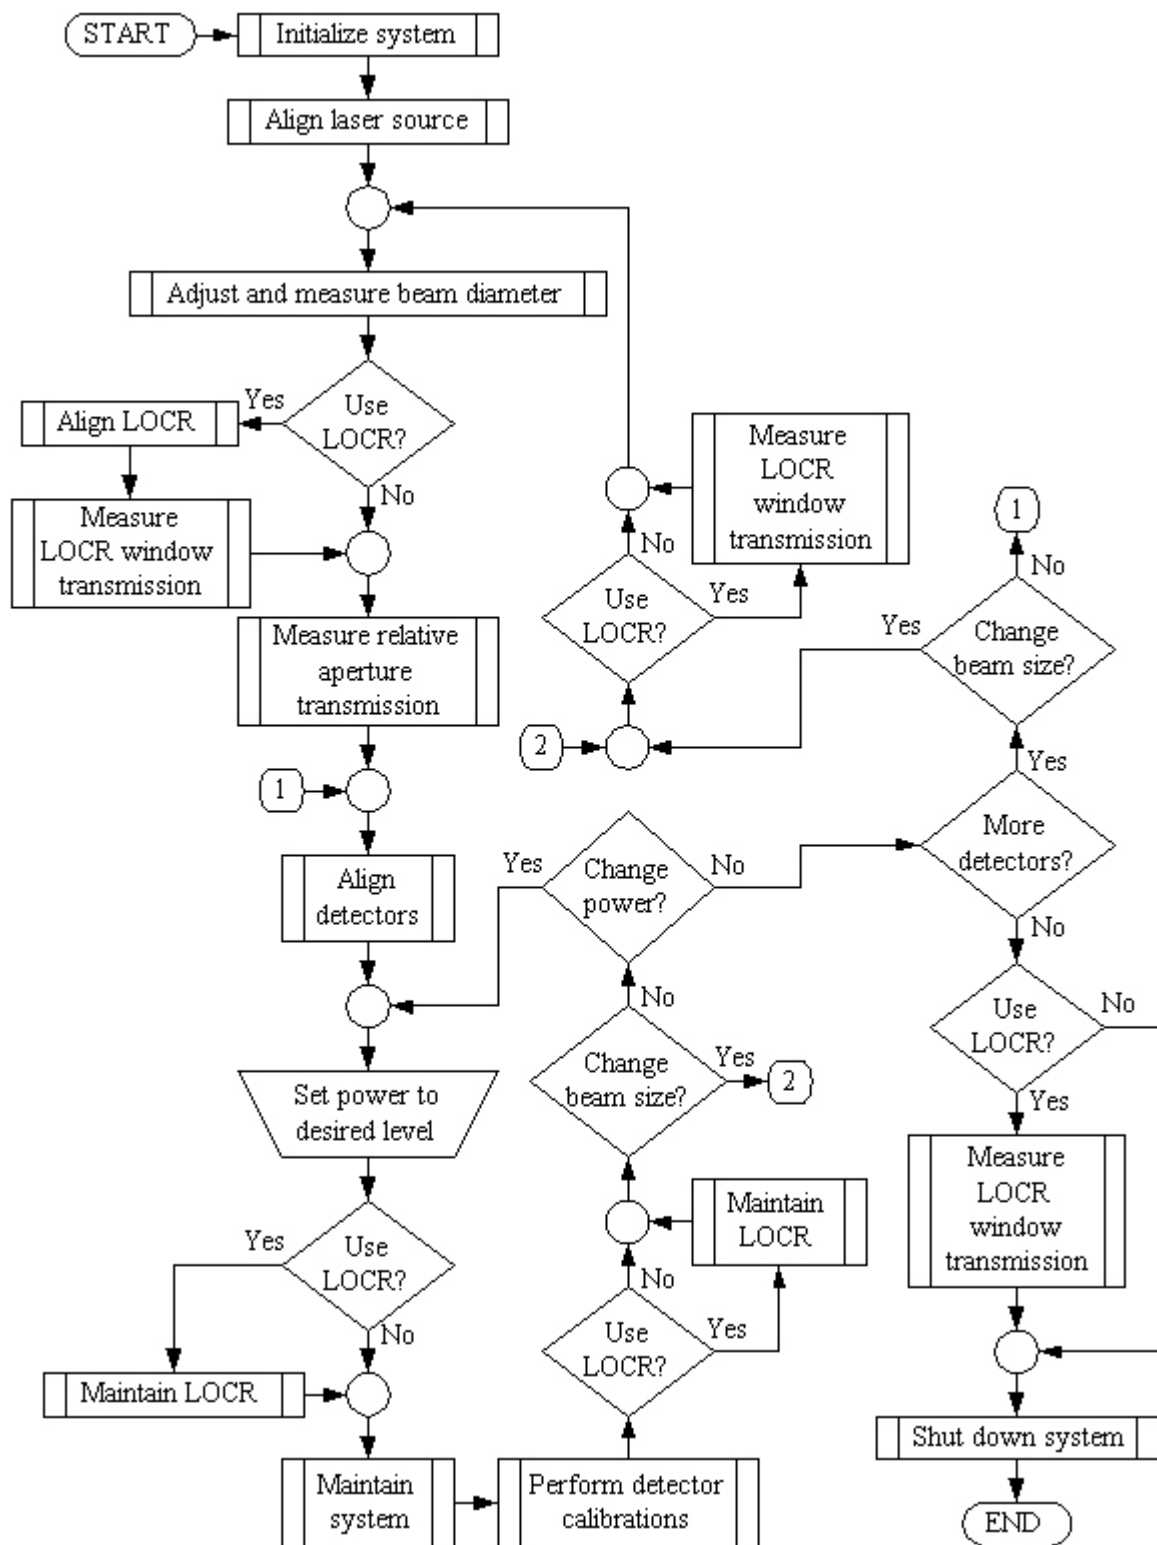


Figure 13. Main calibration process flow chart, for calibration at a single wavelength.

4. CALIBRATION PROCEDURE

Operation of the high-accuracy calibration system is too complex to explain here in detail, but complete operating flowcharts for the high-accuracy calibration system are available in the form of an unpublished document called the High-Accuracy Calibration System Operating Instructions [14]. In the operating instructions, the calibration procedures are described as processes, using flow charts. The flow chart for the main calibration process is shown in Figure 13. The flow chart shows the calibration process for detector calibrations at a single wavelength; the entire process is repeated for additional wavelengths. Here, the operating procedures are broken-up into three fundamental components: operation of the calibration system, operation of the LOCR primary standard, and calculation of the test detector calibration factor.

4.1 Calibration System Operation

4.1.1 Data acquisition software: The calibration system is controlled with software that implements the processes described in the operating instructions. Most of the system is controlled using software implemented on the main process-control computer, which controls the system and acquires the data. Only preliminary data processing is performed during data acquisition, using nominal values for parameters that are measured during the calibration. Final values are unavailable until after post processing is complete.

Customizing the software for a particular test detector is accomplished using the software's graphical interface, which permits real-time editing of the commands sent to the detector, and control of how the detector's output is measured. The software monitors and displays the transactions to help the operator debug the customized instructions. The program traps most of the errors that can occur and has many other built-in checks to help ensure proper operation.

Currently the LOCR is controlled using software provided by its manufacturer, which is executed on a separate computer. The LOCR is treated as a direct-reading detector in the main system software, with a few special extensions that aid in coordinating the two computers and transferring the acquired data between them. LOCR operation is described in section 4.2.

4.1.2 Optical component selection and placement: Component selection is based on the desired wavelength of operation, the particular laser's spectral characteristics, the desired absolute power, and the desired detector-plane beam characteristics. The first requirements determine which laser source should be used, and the choice of laser determines which interference filters (if any) and attenuators are needed. Determining which components are needed to produce the desired detector plane beam characteristics is more difficult. The components of the optical source are shown in Figure 5.

A Gaussian-beam model is used to predict the detector-plane beam size and wavefront based on the selection and placement of the optical components, before the components of the source are actually assembled and aligned. Using the model, different configurations can be considered much more quickly than would be possible if each configuration had to be tested experimentally. Therefore, whether or not the calibration beam desired by the customer can be created using the existing optical components can be quickly determined. Modeling of the Gaussian beam is performed using a ray-tracing analysis that assumes the beam is a perfect Gaussian; a purely geometric analysis is insufficient and can be misleading. In Appendix D, the Gaussian-beam ray-tracing model is described and contrasted with a geometric analysis.

When customers submit a detector for calibration, the minimum information they must provide is the wavelength and power level of the desired calibration. However, they may also specify the desired beam size, wavefront, or divergence. For example, they may request that a Gaussian beam waist with a specific diameter occur at their detector's input, or that the beam diverges or converges as it enters their detector. These parameters are controlled by the selection and location of the components used in the optical system. While all combinations of detector-plane beam size and divergence are not possible, there is a wide range of combinations that can be produced using the existing optical equipment.

The size of the beam in the detector plane is determined by a number of factors, including: the location of the laser source relative to Lens 1, the size and location of the Gaussian beam waist produced by the source; the focal lengths of Lens 1 and 2; the distance between Lens 1 and 2; and the distance from Lens 2 to the detector plane. Optimal collimation, which results in the smallest diameter beam in the far field, is achieved with a specific spacing between Lens 1 and 2. However, the beam can be varied significantly from optimal collimation to obtain different detector-plane beam size, by adjusting the spacing between the lenses.

Initially, the lens spacing that results in optimal collimation is assumed. The lens focal lengths and distance to the laser and detector plane are varied to obtain a detector plane beam that is close to what the customer desires, using the limited set of lens focal lengths available, and laser and detector plane distances that can be accommodated by the system. The detector-plane beam parameters are then fine-tuned by adjusting the spacing between the lenses. The fine-tuning allows fine adjustment of the detector-plane beam size and divergence angle, but they are not independent variables, so one parameter must be specified. For example, suppose that the customer wants the Gaussian beam waist to occur at their detector's entrance aperture. The lens spacing can be tuned to control the location of the beam waist, but the beam's size will also change with the lens spacing. If the resulting beam size is too far from what the customer requires, lenses of different focal lengths must be used, or the laser or detector-plane distance adjusted. Usually the customer is more concerned about the beam's size and the divergence angle is unimportant so long as it is small.

However, the actual beam used in the system is an imperfect Gaussian beam and the optical elements can cause distortion. The result of the imperfection is that the actual beam produced can be slightly different from that predicted by the model. The fine-tuning provided by adjusting the lens spacing is usually sufficient to accommodate any error resulting from the imperfections. In practice, the beams used in the system are good approximations to true Gaussian beams and the lenses are nearly diffraction limited, so the actual beam size and divergence is usually within a few percent of that predicted by theory. To help ensure that the desired beam can be obtained, it is best to choose components and distances such that the desired detector plane beam is well within the limits of the tuning provided by adjusting the lens spacing.

The predicted diameter of the beam waist that occurs between the lenses is used to select the size of the pinhole that is used in the modified spatial filter. A circular pinhole aperture is used with a diameter that is at least twice the theoretical Gaussian beam waist's $1/e^2$ intensity diameter. The relatively large pinhole minimizes the artifacts produced when the Gaussian beam waist is passed through the finite aperture and effectively reduces the scattered power that surrounds the approximately Gaussian beam. It also reduces the system's sensitivity to the laser's pointing stability.

4.1.3 Laser source alignment: Once the component values and distances are selected using the Gaussian beam model, the components are assembled and aligned. Alignment of the laser source consists of placing all the relevant components on the Optical Source Breadboard, shown in Figure 5, and aligning them so that the Gaussian laser beam travels over the center of the rails, at a constant height between components.

In the calibration report, the beam used in the calibration is described in several ways. The beam is described by its $1/e^2$ intensity diameter and its full-width divergence angle, in the detector plane, and by its Gaussian beam waist diameter and location relative to the detector plane. The parameters are determined by direct measurement, or theoretical values are given. The detector-plane beam diameter is always measured; the scanning-slit beam profiler is used to measure the diameter, and the average of the vertical and horizontal width is stated.

The Gaussian beam waist diameter and location are directly measured using the scanning-slit profiler, if a real beam waist occurs at an accessible location. Sometimes the waist cannot be directly measured because it occurs inside the optical source between Lens 2 and the Final Aperture, where it cannot be accessed, or a virtual waist occurs. When the waist cannot be directly measured, theoretical values obtained from the Gaussian beam model are used. Similarly, when the detector plane is in the far-field of the Gaussian beam waist, its divergence is easily measured by performing two diameter measurements a small distance apart. But if the detector plane is very near the beam waist, the divergence angle changes with distance, so instead, the theoretical radius of curvature is used to calculate the divergence angle in the detector plane. Any theoretical values given in the report are described as approximate, while measured parameters are given with their expanded uncertainty.

After the optical components are placed on the rails and aligned as described, the location of Lens 2 is adjusted while the detector plane beam is monitored, until the desired detector plane beam is produced. For example consider that a beam with a particular detector plane beam diameter is desired. In this case, the scanning-slit beam profiler is placed in the detector plane to measure the beam's $1/e^2$ intensity diameter in real-time, as the lens spacing is adjusted as described below while the measured beam size is monitored. However if a particular beam-waist location is desired, the waist location must be determined between each adjustment of the lens spacing.

To adjust the detector plane beam size, the location of Lens 2 is first adjusted for optimal collimation, which is when the far-field beam diameter is minimized. The adjustment is performed by observing the size of the beam a long distance from Lens 2, while the lens is moved. A long distance is achieved by bouncing the transmitted beam back to the operator with a mirror placed near the far wall. Then, the location of Lens 2 is adjusted to achieve the desired detector-plane beam diameter.

The size of the beam in the detector plane is adjusted by varying the location of Lens 2 around the location that results in optimal collimation. To increase the detector plane beam size, Lens 2 is moved closer to Lens 1. To decrease the beam size, Lens 2 is moved farther from Lens 1. The beam diameter varies nearly linearly with lens position over a small range, as described in detail in Appendix D. The detector-plane beam diameter is measured using the scanning slit beam profiler while the lens location is adjusted. If the beam is elliptical, the size is adjusted so that the average of the vertical and horizontal diameters is the desired diameter.

Imperfections in the Gaussian beam cause the beam's actual behavior to differ slightly from the theoretical results. The theoretical analysis usually predicts the beam's minimum detector-plane beam diameter, Gaussian beam-waist location and size, and the far-field divergence angle within about 10 % of the measured results. The difference from the theoretical results is consistent with the beam's quality; the more Gaussian the beam is, the more accurate the theoretical results are. The beam quality is usually limited by what is produced by the laser source itself. Therefore the quality of the laser mirror's alignment (for free-space laser beams), or the beam collimator (for single-mode optical fiber laser sources), are very important and should be optimized to produce a Gaussian beam of the best quality possible.

Currently, a commercial device is used to measure the calibration beam's intensity profile. The commercial device uses a scanning-slit beam diameter measurement technique [15]. The scanning-slit technique accurately measures a Gaussian beam's profile, so is appropriate for the approximately Gaussian beams that are used in the high-accuracy calibration system. The device performs the measurement very quickly, so is especially useful when adjusting the detector-plane beam diameter. The $1/e^2$ intensity diameter that is given in the calibration report is derived from measurements obtained with the device.

A more careful scanning-slit beam profile measurement, performed with customized equipment, can have greater accuracy than that obtained with the commercial scanning-slit beam profiler. More accurate measurements of beam profile can be obtained by measuring the detector's output more accurately and using a precision motion-controlled stage to translate the slit. Such a precision scanning-slit system is currently used in the spatial uniformity scanning system [10]; the precision system provides reproducible beam diameter measurements with a standard deviation of typically 0.02 %. In contrast, the commercial beam profiling system's measurements based on a scanning slit are relatively noisy, with a typical standard deviation of 0.4 %. The speed of the commercial device is helpful while the beam size is initially adjusted, but a precision width measurement is more desirable for the calibration report. Precision scanning-slit or pinhole beam profiling measurements may be performed using the precision x-y sub-stage, described in Section 3.3.

Once the laser source is aligned, the laser beam's height is controlled by adjusting the height of the Optical Source Breadboard; otherwise the beam's direction is fixed and the detectors are aligned to the beam's position. The height adjustment of the source breadboard is necessary to accommodate the LOCR, which has a fixed height. All other devices are aligned to the resulting beam height. Horizontal positioning is aided by the motion-controlled detector stage.

After the beam is aligned and its height is adjusted, its location in the detector plane is fixed and everything else must be aligned to that location. The beam power can subsequently be changed by adjusting the LPC and attenuators, without changing the location of the beam in the detector plane. Therefore, the beam power can be adjusted without requiring detector realignment or beam parameter re-measurement. Changing the beam's size necessitates remeasurement of the beam parameters that affect measurement uncertainty, such as the relative aperture transmittance, LOCR window transmittance and LOCR spatial uniformity. Similarly, changing the laser wavelength requires a complete realignment of the laser source and detectors, and remeasurement of the beam parameters.

4.1.4 Detector alignment: After the optical source is aligned and the desired detector plane beam is produced, the detectors are aligned. First the LOCR is aligned, but because the LOCR's height is not adjustable, the height of the optical source breadboard is adjusted instead. LOCR alignment is described in detail in Section 4.2.2. After the LOCR is aligned, the test detectors are aligned.

To help describe the alignment procedure, the naming convention shown in Figure 14 is used. Usually the detectors are mounted using a post and holder, as depicted in the figure. With such a mount, no pitch or roll adjustments are available; the angles are set by the physical characteristics of the mount. However, more complicated mounting platforms can have pitch and roll angle adjustments.

Unless the customer demands a specific alignment procedure, the test detectors are aligned so that the beam's incidence angle is as close to normal as possible relative to the plane of the detector's entrance aperture. The exception to this rule is when a strong back-reflection from the detector re-enters the optical system, when the detector is aligned at normal incidence. In this case, the detector's yaw angle is changed enough to prevent the reflection from re-entering the optical system.

The beam's incidence angle is adjusted relative to the detector's physical entrance aperture. To determine the beam's incidence relative to the aperture, the operator holds a glass flat up to the aperture, so that a strong specular back reflection is produced. The reflection from the glass flat is compared to the incident beam to adjust the incidence angle. The reflected beam will fall on top of the incident beam when the incident beam is at normal incidence relative to the aperture.

Before the angular alignment begins, the detector mount's location is adjusted in the longitudinal direction so that the detector's limiting aperture is in the detector plane. If the location of the detector's limiting aperture is unclear or the detector has no limiting aperture, the detector's entrance aperture is positioned in the detector plane instead. Then the detector's pitch angle is adjusted, if available. The pitch angle is adjusted so that the beam reflected from the glass flat is in the same horizontal plane as the incident beam.




Figure 14. Rotational and translational motions used during detector alignment.

Next the detector's height is adjusted, using a vertical translation. The vertical translation is usually provided by sliding the post in its holder. Once the height is adjusted, a post collar is locked into position so that the post can be rotated without changing its height. The height is adjusted so that the beam passes through the vertical center of the detector's entrance aperture. The aperture's center can be located by eye using a ruler, or using the alignment aid located on the detector's aperture cover, if any. A more precise alignment can be obtained by fitting the detector with an aperture cover that has a pinhole aperture in its center. Then the vertical height is adjusted until the power passing through the pinhole aperture is peaked. This technique is most commonly used with detectors intended for use with optical fibers, then the optical fiber adapter itself is used as the pinhole aperture, to ensure that the collimated beam strikes the detector in the same location that a fiber-coupled beam would.

After the detector's height is adjusted, its yaw angle is adjusted. Usually, yaw is adjusted by rotating the post in its holder; the post collar prevents any change in the detector's height when the post is rotated. As with the pitch angle, the yaw angle is adjusted by observing the reflection from the glass flat that is held against the detector's entrance aperture. The yaw angle is adjusted so that the back reflection is in the same vertical plane as the incident beam.

Then the beam is centered in the detector's aperture horizontally. The horizontal translation is performed with the motion-controlled detector stage, using the y-axis of the joystick. As with the vertical centering, the alignment is performed visually, using a pinhole aperture, or with whatever special technique the customer requests.

Lastly, the detector's roll angle is adjustment, if any, is performed. Usually, a roll adjustment is used only with detectors that are sensitive to the polarization angle and have radial symmetry. With such detectors, the roll angle can be adjusted without affecting the beam's centering or incidence angle. If the roll adjustment does affect the other parameters, the roll angle should be adjusted first instead of last. Usually the detector's roll angle is defined by the physical characteristics of its mount, so is not adjustable.

A conservative uncertainty estimate of the error in the detector's alignment is assessed and stated in the calibration report. Typical upper-bounds are 1 mm for visual centering using a ruler, 0.5 mm for visual centering using an aperture cross-hair or other targeting device, and 0.1 mm for centering by peaking the power transmitted through a pinhole aperture mounted on the detector. The alignment uncertainty is used in conjunction with the detector's spatial uniformity to determine the uncertainty due to an imperfect alignment.

Similarly, a bound for the error in the laser beam's incidence angle is also assessed and stated. The difference between the laser beam's actual incidence angle and normal incidence is determined by holding the glass flat to the detector's entrance aperture, locating the back reflection cast by the glass, and determining the angle between the back reflection and the incident beam. The difference between the actual incidence angle and normal incidence is half the angle between the incident and reflected beams. The measured angle is combined with its measurement uncertainty to give an upper bound for the difference between the beam's actual incidence angle and normal incidence, which is typically a few milliradians. The direction of the angular alignment error is not specified.

4.1.5 Relative aperture transmittance measurement: After the optical source is aligned, and usually before the test detectors are aligned, the relative aperture transmittance is measured. The measurements are relevant only for a particular laser beam profile, so must be performed whenever the beam's profile is changed, such as when the optical system is adjusted for a different beam size. Once the optical source is adjusted to deliver the desired beam profile, the relative transmittance is assumed to remain stable for the duration of the calibration, so the measurement is performed only once during the calibration. Also, it is assumed that the beam profile and hence the relative transmittance does not change when the beam power is adjusted, so the power level can be adjusted without requiring another set of relative transmittance measurements. A power level of 1 mW is usually used during transmittance measurement. The relative aperture transmittance measurement process is discussed in detail in Appendix E, so is described only briefly here.

To determine the relative aperture transmittance for a particular test detector, the shape and size of its apparent aperture is needed. When the detector's apparent aperture must be measured, because the detector lacks a limiting aperture, a figure showing the measured apparent aperture is provided in the calibration report. The detector's apparent aperture is measured with a spatial uniformity scan. When a physical limiting aperture is assumed, the assumed shape of the aperture is described in the report text; such apertures are usually circular. A spatial uniformity scan of a detector with a physical limiting aperture may or may not be performed, depending on the customer's handling specifications. If the customer allows the detector to be cleaned (by blowing with pressurized air, for example), uniformity scans are performed to determine whether the detector is dirty, and if so, whether subsequent cleaning was effective. In this case, a figure showing the remaining contamination, if any, may also be included in the report. But a formal report of spatial uniformity is not provided as part of the high-accuracy responsivity calibration.

The relative aperture transmittance is usually measured using a laboratory-standard detector and fixed circular apertures. Whether or not the test detector itself is used to measure relative aperture transmittance is dependent on its apparent aperture's shape and size. If the test detector does not have a physical limiting aperture, or if its aperture is larger than the largest fixed circular aperture, the test detector itself is used to measure relative aperture transmittance.

The laboratory-standard detector that is used to measure relative aperture transmittance has an apparent aperture that is at least as large as the largest fixed circular aperture, and has good spatial uniformity within its aperture. A large single-element Si or Ge detector is used as the standard detector; the Si

detector is used with visible wavelengths, and the Ge detector is used with near-infrared wavelengths. The relative transmittance of the fixed circular apertures is measured using the standard detector. The detector is angled slightly so that its back reflection does not re-enter the optical system. When placed in front of the detector, the apertures are very close to the detector's surface, so that no light that is transmitted through the aperture is diffracted away from the detector. The apertures are painted with a flat-black paint to help absorb any light that is reflected off the detector, so that it is not reflected back to the detector.

First, the standard detector is aligned so that its entrance aperture is slightly behind the detector plane and the beam is centered in its aperture. Then each fixed circular aperture is in turn placed in front of the detector's aperture, in the detector plane, so that the beam is centered in the aperture. One aperture, usually the largest, a circular aperture of 1 cm diameter, is used as the reference; the transmittance of each of the other apertures is measured relative to the transmittance of the reference.

Each fixed aperture's transmittance relative to that of the reference aperture is measured. First the reference aperture is placed in front of the standard detector and the power received by the detector is recorded. Then the reference aperture is replaced by the aperture under test and the power received by the detector is recorded. The transmittance of the test aperture relative to that of the reference aperture is then the power recorded when the test aperture was used divided by the power recorded when the reference aperture was used. The absolute beam power is assumed to be stable or vary linearly during the measurement. To reduce the affect of any power drift, the power during the test aperture measurement is interpolated in time using bracketing measurements made with the reference aperture, and multiple measurements are averaged.

If all of the detectors involved in the calibration have true limiting apertures that are within the range of the fixed circular apertures, no additional measurements are needed, even if a limiting aperture is non-circular. For example, the LOCR has a physical limiting aperture that is elliptical, with dimensions 8.5 by 9.5 mm. The transmittance of its aperture is therefore somewhere between the transmittances of the 8 mm and 10 mm diameter circular apertures, so can be estimated using the transmittance data for the 8 and 10 mm diameter fixed apertures. Basically, the transmittance estimate for the elliptical aperture is the average of the transmittance measurements for the 8 and 10 mm diameter circular apertures, with an uncertainty that includes the change in transmittance between the 8 and 10 mm apertures, as described in the appendix. In general, the circular apertures used in the estimate are (a) the largest centered circular aperture that is completely enclosed within the detector's limiting aperture, and (b) the smallest centered circular aperture that fully encloses the detector's limiting aperture.

Ultimately, the relative transmittance between the LOCR primary standard and each test detector being calibrated is needed. Continuing the example, consider that a test detector with a circular entrance aperture of 5 mm diameter is being calibrated. As described above, the transmittance of the LOCR's elliptical aperture, relative to that of the reference aperture, is the average of the 8 and 10 mm diameter circular apertures, relative to the reference aperture. The transmittance of the test detector's aperture relative to that of the reference aperture is therefore the relative transmittance of the 5 mm diameter circular aperture. Then the transmittance of the LOCR's aperture relative to that of the test detector's is simply the transmittance of the LOCR's aperture relative to the reference, divided by the transmittance of the 5 mm diameter aperture relative to the reference.

The test detector is handled differently when it lacks a limiting aperture, or when the test detector's aperture is larger than the largest fixed circular aperture available. In these cases, the largest centered, circular aperture that is fully enclosed by the test detector's apparent aperture is used, along with the test detector itself, to determine the detector's relative transmittance. For example, the apparent aperture

shown in Figure 6 fully encloses a centered circular aperture with a diameter of 7.8 mm. Since circular apertures of fractional sizes are not used, the circular aperture of 7 mm diameter is the largest that is fully enclosed by the detector's apparent aperture. Similarly, a test detector with a very large aperture fully encloses the largest fixed circular aperture. The test detector itself is used to measure the transmittance of its aperture relative to that of the centered circular aperture, by measuring the beam power when its aperture is unobstructed, and dividing by the beam power measured when the centered circular aperture is placed in front of the detector. This results in the transmittance of the detector's aperture, relative to the enclosed circular aperture. To determine the transmittance of the detector's aperture relative to the transmittance of the reference aperture, the transmittance relative to the enclosed aperture is multiplied by the transmittance of the enclosed aperture relative to the reference aperture. Then the test detector's aperture transmittance relative to the standard's is calculated as described above. Like the measurements performed with the laboratory standard and fixed apertures, the measurement of the transmittance of the test detector's aperture relative to the enclosed circular aperture's is performed using interpolation in time, and multiple measurements are averaged.

In addition to the measurements described here, there are other measurements performed to determine the uncertainty in the aperture transmittance measurement itself. And the measured uncertainties are propagated as the different transmittance measurements are combined to obtain the final result. The measurements and calculations are fully described in Appendix E.

4.1.6 Calibration system maintenance: Several of the components of the high-accuracy calibration system require periodic maintenance. Maintenance of the system components that are independent of the LOCR are discussed here. Maintenance of the LOCR primary standard and equipment used exclusively with it is discussed in section 4.2.6.

The calibration system is powered through an uninterruptible power supply (UPS) that provides brown-out protection as well as electrical isolation and noise filtering. The UPS has a lifetime that is limited primarily by its internal battery, so the battery's health must be monitored. The device performs a monthly self-check and emits an audible and visible alarm if a problem is detected. The details of the self-check and other status can be read from the UPS through its RS-232 serial interface. When problems are detected, the UPS must be repaired or replaced.

The HEPA filter that filters the air inside the lab table enclosure also has a limited lifetime. The HEPA filter element's lifetime is determined in part by the amount of dirt present in the environment. Since the filter cannot be inspected without disassembly of its enclosure, it is usually best to simply replace the filter element periodically. The original filter was replaced after about five years of continuous operation; at that time, the once-white filter element had a brown color. Since HEPA filter elements can theoretically be operated until they are essentially black and impenetrable, the filter could probably have been used longer. However, the rate of air flow through the filter decreases as the filter becomes clogged, so a periodic replacement at the rate of once every five years is preferable.

The calibration system's precision DMM requires periodic maintenance in the form of electrical calibration. The DMM is the primary standard for all electrical measurements performed in the lab, so the uncertainty of its measurements limits many of the other uncertainties in the calibration. The accuracy of the DMM calibration is affected by temperature and time. Currently, the manufacture's uncertainty specification for a time period of 90 days and for a temperature range of the DMM's calibration temperature $\pm 5^\circ \text{C}$ is used in the uncertainty analysis. Therefore, the DMM must be electrically calibrated within 90 days of when it is used to perform detector calibrations, and the temperature during the calibration must be within $\pm 5^\circ \text{C}$ of the temperature present during its electrical calibration. Currently, the DMM electrical calibration is performed in-house, and a matrix of measured

voltages and voltage error is provided. The measured voltage errors are corrected in post-processing, as described in Appendix F.

The software used to control the calibration system will require maintenance in the form of adding new features and debugging, in addition to migrating the system to newer hardware as computer technology improves. To help in this maintenance, the computer programs are written using standard structured techniques in a simple Basic language, so should be understandable by any computer programmer. The user interface is graphical and intuitive, so an experienced laboratory technician should be able to operate the system, after some training or experimentation. The software stores all the acquired data in standard spreadsheet format, along with a description of the measurement performed. Most of the mathematical calculations are performed on the spreadsheet, so that the analysis performed is documented.

4.1.7 Quality assurance checks: Throughout the calibration processes, several quality assurance checks are performed. The checks are performed to identify problems before they become evident as bad calibration results. For example, the optical source's power stability is verified before any calibrations are performed. The processes used in the calibration are documented in the High-Accuracy Calibration System Operating Instructions [14]; strict adherence to the processes prevents missed steps that can add unaccounted uncertainty to the calibration results.

In one of the more visible quality-assurance checks, a trusted transfer standard or "check" standard is usually calibrated alongside every test detector calibration. If the check standard's calibration changes, a problem with the calibration system is indicated. Since the calibration system has the ability to calibrate more than one detector at a time, calibrating the check standard alongside the customer's test detector consumes little additional time and resources.

In addition to the check standard, there are many other checks that occur during the calibration because previously obtained results should be reproduced. If unusual results are obtained, a potential problem with the calibration system is indicated. These system consistency checks are in some ways more useful than the result of the check standard, because they can indicate a problem earlier in the calibration process, while the entire calibration must be completed before the check standard's result is known.

One example of a system consistency check is the electrical calibrations performed. In these checks, a result consistent with previously obtained results is expected; if unusual results are obtained, a problem with the system is indicated. For example, the absolute accuracy of the LOCR's electronics is measured during each calibration under the same environmental conditions that are present during the calibration. The calibration system's DMM is the standard for all electrical measurements, it is electrically calibrated within 90 days of the detector calibration. The electrical calibration factors for both devices should be consistent, unless there is a reason for a change, such as a relatively large change in the environment or a long time span between calibrations. Other similar consistency checks are performed during the calibration, such as when the optical system's stability is assessed, when the beam's relative aperture transmittance is measured, and when the LOCR's uncertainties are assessed.

In perhaps the most visible and reassuring check, the system's absolute performance is empirically tested through comparison with other high-accuracy devices, usually other cryogenic radiometers. To date, transfer standards have been used to compare the LOCR to the HACR cryogenic radiometer at 488, 515, and 633 nm [16]; and the C-series calorimeter at 633 and 1550 nm [6]. Such comparisons will be performed periodically.

When previously unaccounted uncertainties are identified, they are added to the error budget as soon as possible. For example, a recent publication [17] showed that the spectral response of the optical receiver

in a device similar to the LOCR was not flat. Theoretical analysis based on the reflection characteristics of the paint and the geometry of the receiver concluded that the spectral response should be essentially flat from visible to near infrared wavelengths, so the optical receiver's absorption was assumed to be spectrally flat. But the recently measured results show that the spectral response is not as flat. Until proven otherwise (the published variation may be the result of unaccounted measurement uncertainty), or additional measurements are available, the published results will be used to reassess the uncertainty due to spectral response in the LOCR. Similarly, after discovery that the calibration system's beam contains some non-Gaussian components, the relative aperture transmittance correction was added to better assess the resulting uncertainty. And most recently, significant spatial non-uniformity in the LOCR's optical absorber was measured and accounted for. The uncertainty analysis is described in Section 5. While including these uncertainties may make the system's overall calibration uncertainty somewhat larger than that advertised for other laboratories, it better estimates the true uncertainty in the calibration results.

In part because of the quality assurance checks, the final calibration factor and uncertainty cannot be determined until after the calibrations are completed and all the measured corrections and uncertainties are tallied. Fixed values for measured uncertainties that prove to be very stable with time may eventually be used after a sufficiently large calibration history is obtained, but doing so would bypass the quality assurance check performed when the parameter is remeasured for each calibration. Currently, using typical values for the measured uncertainties gives a preliminary result with an uncertainty that is usually less than 0.1 %, which is sufficient to help determine whether the system is functioning properly while the calibrations are being performed.

To help ensure the accuracy of the calibration result, conservative uncertainty estimates are used when there is doubt. For example, the LOCR window transmittance measurement is performed extremely conservatively. Instead of measuring the transmittance at one location on the window one time, a matrix of locations both before and after the calibrations is measured, to account for uncertainty in the window's alignment and changes in its transmittance that can occur during the calibration. If the window transmittance changes significantly, the uncertainty of the calibration results will be increased correspondingly, so the calibrations can be delivered with an uncertainty that is larger than expected, or the calibrations can be repeated.

4.2 LOCR Operation

Operation of the LOCR primary standard is described in great detail in the manufacturer's operating instructions [13] and in the operating instructions for the high-accuracy calibration system [14], so is only summarized here. Operation of the LOCR must be performed with great care, because mistakes or missed steps can result in permanent damage to the device or injury to the operator.

4.2.1 Vacuum and cryogenic systems: The cryogenic liquids needed for LOCR operation present several safety hazards. Oxygen (O_2) depletion is probably the greatest hazard for the operator, so the O_2 level must be monitored. The O_2 concentration of the lab air is monitored with a commercial O_2 sensor that displays the concentration and emits audible and visible alarms if the concentration drops to hazardously low levels. The O_2 depletion occurs because whenever cryogenic liquids are exposed to the relatively warm atmosphere and equipment, the liquid boils and becomes a gas. The large amounts of N_2 gas released can easily displace enough of the O_2 in the room to cause problems for the human operator unless the air in the room is constantly replaced in an air exchange. In the high-accuracy calibration system, the room air exchange is provided by a fan that blows the room air through a vent in the wall to the outside of the building. To aid in the removal of the unwanted gasses from the laboratory, the air that is exhausted is drawn from a vent located directly above the LOCR. Latex hoses are used to help direct

the gas to the vent, where appropriate. The ventilation fan's speed is controlled from inside the laboratory so that the air exchange rate can be increased when needed and reduced when not needed.

The He gas causes less displacement of O₂, because there is usually much less He gas released and the He gas tends to escape from the room relatively quickly. However, even a small concentration of He in the room's air can change the air's index of refraction. The air's index is calculated from measurements of the air pressure, temperature, and relative humidity; and the calculation assumes typical values for concentrations of the gasses found in Earth's atmosphere. The presence of excessive He or N₂ gas in the room air could therefore invalidate the calculation of the index of refraction. Fortunately, the location of the vent directly above the LOCR helps eliminate most of the gasses released and the room air exchange reduces the remaining excesses to negligible levels.

The cryogenic liquids present hazards for the operator in addition to the O₂ depletion hazard. For example, the operator can be burned or temporarily blinded if any cryogenic liquid contacts the skin or eyes. To prevent these potential injuries, insulated cryogenic gloves and goggles are worn whenever the cryogenic liquids are handled. Quality cryogenic gloves are used, which consist of a waterproof shell and an insulated liner. The shell allows the gloves to withstand immersion in the cryogenic liquid for a short time, while the insulated liner protects the skin from freezing. To protect the operator's arms, a long sleeve shirt or laboratory jacket is worn when handling the cryogenic liquids. Safety goggles that completely cover the operator's eyes are worn to prevent any splashed liquid from reaching the eyes.

The equipment itself can also be damaged by the cryogenic liquids. When the liquid nitrogen first contacts the relatively warm metal, it boils furiously and small amounts of liquid can be sprayed around the area. Also when the LOCR's liquid-nitrogen reservoir is filled, there is usually some overfill that can spill onto the detector breadboard. To protect the test detectors and other equipment located on detector breadboard, a thin metal shield is used to deflect any sprayed liquid away from the equipment. The liquid helium is a lesser problem because it boils so quickly that liquid helium is not observed outside of the cryogenic containers.

The greatest danger to the equipment results from the tendency for ice to form on the cold surfaces. The cryogenic liquids produce temperatures that are cold enough to condense atmospheric water vapor (and other atmospheric constituents). The worst possible scenario is for an ice plug to form in the LOCR's cryogenic reservoir fill ports. If the ports are clogged, the pressure inside the reservoirs will buildup until the ice plug is ejected, or the reservoir itself ruptures. There are no pressure relief valves on the LOCR's cryogen reservoirs. However, the LOCR's cryostat was designed to minimize the potential for ice plugs to form.

As shown in Figure 8, the LOCR's liquid-nitrogen reservoir shields the liquid-helium reservoir from the warm sides of the cryostat. Because of the reservoir's exposure to the warm cryostat, the liquid-nitrogen boils vigorously, even after the reservoir reaches its steady-state temperature. The rapid rate of boiling produces a steady stream of N₂ gas that helps to keep outside air from entering the reservoir and freezing. Additionally, the LOCR has two fill ports connected to its liquid-nitrogen reservoir, which provide redundant paths for venting; if one port freezes, the rate of gas flow through the second port increases, thereby lowering the risk of ice forming in the second port. In practice, latex tubing is placed in the liquid-nitrogen reservoir's fill ports to direct the escaping gas to the vent hood. To date, no icing inside the latex tubes that vent the N₂ gas has been observed, though ice does accumulate on the outside of the tubing. However to ensure that a problem does not occur, the N₂ gas vent hoses are inspected at least twice a day, when the liquid-nitrogen reservoir is refilled.

The liquid-helium reservoir is significantly different from the liquid-nitrogen reservoir and presents a greater danger. The primary difference is that there is only one fill port on the liquid-helium reservoir and its port is vacuum insulated so it is very cold. Also since the liquid-helium reservoir is well shielded by the liquid-nitrogen reservoir and its heat shields, relatively little He gas is boiled off once the reservoir reaches the temperature of liquid helium, so there is little flow of gas through the fill port. To prevent ice from forming inside the liquid-helium fill port, a device called a vapor barrier is inserted into the port. The vapor barrier consists of three metal barriers that act as heat shields between the cold liquid helium and the room-temperature atmosphere, as well as physical barriers to gas flow. The vapor barrier fits underneath the liquid-helium reservoir's cap, but the barrier must be removed during cryogenic liquid fills. After filling the reservoir with liquid helium, the cap must be removed, the vapor barrier inserted, and the cap replaced. Failure to insert the vapor barrier after a fill would very likely cause an ice plug to form.

The He gas produced by the boiling liquid vents through a horizontal fitting in the liquid-helium reservoir's cap, located above the vapor barrier. The fitting is designed so that a vacuum pump can be attached and the pressure in the reservoir lowered, to lower the temperature of the liquid helium. In an experiment, a length of latex hose was attached to the fitting to direct the escaping He gas to the vent hood above, as is done with the N₂ gas. While this technique is successful with the cold N₂ gas which is heavier than air, when used with the lighter-than-air He gas, the formation of an ice plug on the vapor barrier is enhanced by the hose. Therefore a latex exhaust hose is not used on the He gas vent. However even without a hose attached, the potential for formation of ice on the vapor barrier still exists, so the liquid-helium fill port is inspected for the presence of ice at least twice a day. Also, ice formation in the liquid-helium reservoir can be inferred by monitoring the rate of liquid-helium consumption; if ice is being formed, the liquid helium will likely be consumed at a greatly accelerated rate and the liquid-helium temperature will increase if the pressure in the reservoir increases.

The storage container for liquid helium also requires special attention. Like the LOCR's He reservoir, the storage container also has only one fill port, which presents a similar icing hazard. The container is safe when its fill port is properly closed and capped, but the port must be opened and uncapped to insert the liquid-helium transfer line. While the port is open there is a danger of ice buildup and if an ice plug forms, the pressure in the container can rise to dangerous levels. Fortunately the storage container has several safety features which should prevent an explosion even if an ice plug does form, but the container may still be damaged and the backup pressure reliefs should not be relied on. Instead any ice plug that forms should be removed immediately. A long copper rod about 1 cm in diameter is used to remove the ice plug; the rod is inserted into the fill port and used to push the ice plug into the container. The best policy is to ensure that the container is properly capped when not in use and is uncapped for only a short time, when necessary.

In addition to the icing that can occur inside the LOCR's liquid cryogen fill ports, there is also an issue with ice forming on the outside of the cryostat. The top of the cryostat tends to accumulate ice during liquid-nitrogen fills. If the cryostat is allowed to freeze, its o-ring vacuum seal could also freeze and would then likely rupture. The resulting inrush of air and subsequent ice formation would likely destroy the LOCR, so care must be taken to ensure that o-ring seal does not freeze. To prevent such icing, a heater designed to prevent water pipes from freezing is wrapped around the top of the cryostat. The heater has an internal thermostat so should operate automatically, but it should be inspected to ensure that it is operating correctly. As a backup, a commercial hot air gun is used to manually remove any ice buildup. The hot air gun is also usually needed to warm the liquid-helium reservoir's cap so that it can be removed for fills, since the cap tends to become stuck after being exposed to the liquid-helium vapor for a time.

Care must also be exercised when using the vacuum system, or damage can result. For example, if a vacuum seal or a valve is opened at the wrong time, a inrush of air can occur, which would likely destroy the LOCR's fragile optical receiver. Also, the cleanliness of the turbo pump and its associated plumbing must be maintained. To help prevent contamination, the vacuum system components are always handled with gloved hands, and the pump is always vented with dry nitrogen gas after use. These and other maintenance functions are discussed in greater detail in Section 4.2.6.

The first step in using the LOCR is to evacuate the cryostat and window assembly, using the turbomolecular pump station. When the pressure in the LOCR's cryostat drops below about 0.005 Pa, the ion pump is started. The ion pump will be running hot at this pressure, with an unusually high operating current. The heat will make the ion pump outgas, which will drive the pressure back up, which will make the pump run even hotter, etc. in a positive feedback loop which will eventually result in destruction of the ion pump or its power supply.

The ion pump power supply has a safety feature that shuts the pump off if it draws more current than a preset limit. Unfortunately, the safety feature cannot be used while the pump is starting, because during startup it draws much more current than the preset limit, so the safety feature must be disabled during startup. After the pump is running with a stable pressure, the current will drop low enough to reenable the safety feature. Then the safety feature is active for the duration of LOCR operation.

To prevent damaging the ion pump while the safety feature is disabled, the operator monitors the pressure, and if the pressure rises to about 0.013 Pa, the ion pump is shut off and allowed to cool, while the turbo pump removes the gasses that were released. When the pressure is reduced to about 0.005 Pa again, the ion pump is restarted and the process is repeated. This cycle is typically repeated a few times, until the ion pump can maintain the pressure without outgassing. Once the ion pump is running with a stable pressure, the valve to the turbo pump is closed, the turbo pump is removed, and the LOCR's liquid-nitrogen and liquid-helium reservoirs are filled with liquid nitrogen. Introduction of the cryogenic liquids will further reduce the pressure in the cryostat and take the load off the ion pump.

The liquid nitrogen is transferred from the storage container to the LOCR using a flexible, insulated transfer line, described in Section 3.1.4. The liquid-nitrogen reservoir is filled first. Before filling the reservoir, the latex tubing is removed from the liquid-nitrogen fill ports. Then the liquid-nitrogen phase separator at the end of the flexible liquid-nitrogen transfer line is inserted into one of the ports. A special clear plastic pipe is placed over the unused fill port; the pipe helps direct the escaping gas to the vent hood, while allowing visual inspection of the port. The flow of liquid-nitrogen is then started, the valve on the liquid-nitrogen storage container is used to control the rate of liquid flow. When liquid is observed overflowing from the exhaust port, the reservoir is full so the valve is closed to stop the flow. Then the pipe and phase separator are removed and the latex tubing replaced.

The LOCR's liquid-helium reservoir is pre-cooled using liquid nitrogen, so that less liquid helium is consumed during the liquid-helium fill. First the reservoir cap and vapor barrier are removed. Then an extension tube is inserted into the liquid-helium reservoir and the liquid-nitrogen phase separator is inserted into the tube. The extension tube almost reaches the bottom of the reservoir, it is held off the bottom by a metal clip which hooks onto the top of the fill port. Then the flow of liquid nitrogen is begun and continues until excess liquid is observed flowing from the top of the fill port. When the fill is completed, the vapor phase separator and extension tube are removed from the reservoir. At this time, the reservoir cap is not reinstalled, so that the liquid nitrogen in the reservoir can be observed visually.

When the liquid nitrogen is first introduced to the LOCR's warm cryogen reservoirs, the liquid will boil furiously. After some period of time, about 30 minutes or so, the LOCR's reservoirs are pre-cooled to

liquid-nitrogen temperature and the liquid nitrogen's rate of boiling will slow. When fully precooled, the liquid nitrogen inside the well-insulated liquid-helium reservoir will practically stop boiling entirely. Then the liquid nitrogen can be blown out of the liquid-helium reservoir, so that the liquid helium can be introduced.

To remove the liquid nitrogen from the liquid-helium reservoir, the liquid nitrogen is blown out of the liquid-helium reservoir and into the liquid-nitrogen reservoir. A special liquid-nitrogen blow pipe is used to remove the liquid. To use the blow tube, the reservoir cap is installed on the liquid-helium reservoir, without the vapor barrier in place. Then the blow pipe is inserted into the cap, until it contacts the bottom of the reservoir. The blow pipe's end is cut at a slight angle and notched so that there is a opening for the liquid to flow through. On the top of the blow pipe is a short length of latex tubing, which is inserted into one of the liquid-nitrogen reservoir's fill ports. Another length of latex tubing connects the vent on the liquid-helium reservoir's cap to the calibration system's gas valves, shown in Figure 4, which are used to control the flow of gas. Then gaseous N_2 is blown into the reservoir through the tube, and the pressure drives the liquid nitrogen up the blow pipe and into the liquid-nitrogen reservoir. Dry N_2 gas is usually used to blow out the liquid nitrogen, but He gas could also be used. It is very important that all the liquid nitrogen be removed, or excessive amounts of liquid helium will be consumed as the remaining liquid nitrogen is frozen solid. To ensure that all the liquid nitrogen is removed, a felt pad is inserted into the liquid-helium reservoir to physically mop up any liquid remaining on the bottom of the reservoir. Then the LOCR's liquid-helium reservoir can be filled with liquid helium. It is important that these steps be performed quickly, so that the liquid-helium reservoir does not warm up again after the liquid nitrogen is removed. There should be enough room in the liquid-nitrogen reservoir to hold all of the liquid that is blown out of the liquid-helium reservoir; if not, the excess is discarded.

The liquid helium is transferred to the LOCR using a rigid, vacuum-insulated transfer line. To control the rate of flow of the liquid-helium through the transfer line, the liquid-helium storage container is pressurized with He gas. A special adapter is used to connect the calibration system's gas valves to the liquid-helium container and the system's gas valves are used to control the flow of gas. At first, the warmth of the liquid-helium transfer line will vaporize the liquid-helium inside the line, so only He gas will flow. The transfer line is not inserted into the LOCR until liquid begins to flow through the transfer line, which may take a few minutes. Then the transfer line is inserted into the LOCR's liquid-helium reservoir and the reservoir is filled. The amount of liquid helium consumed in the filling process is directly dependent on the operator's skill, a sloppy fill will consume much more liquid helium than a properly performed fill. The most efficient transfer occurs with a certain rate of flow of liquid helium through the transfer line; either too fast or too slow a flow can result in excessive consumption of liquid helium. So the operator must carefully monitor the rate of liquid-helium flow, by observing the plume that occurs at the liquid-helium reservoir's vent, and control the flow rate by the application of gas pressure. The fullness of the liquid-helium reservoir can be determined by observing the exhaust plume, or with the liquid-helium level sensor that is built into the reservoir's cap. The level sensor provides a direct reading of the amount of liquid-helium in the reservoir, but an experienced operator can identify when liquid is being ejected from the vent port by the look of the exhaust plume, which becomes denser and more feathery when containing liquid helium. After the reservoir is filled, the liquid-helium transfer line and reservoir cap are removed. Then the vapor barrier is inserted into the cap and the cap is slowly lowered onto the reservoir's fill port.

Once the liquid-helium reservoir is filled with liquid helium, the LOCR's internal temperature will quickly drop to near liquid-helium temperature and the LOCR's temperature-control electronics will begin to function. The electronics will not function at higher temperatures, primarily because the wire leads do not become superconducting until their temperature is well below liquid-nitrogen temperature.

4.2.2 Alignment and spatial uniformity measurement: The first step in aligning the LOCR is to center the laser beam in its optical receiver. The LOCR is first coarsely aligned so that at least part of the laser beam enters the optical receiver. The LOCR alignment target can be used for the coarse alignment, if the target was previously aligned as described below, otherwise the LOCR is coarsely aligned. After at least part of the beam enters the LOCR's receiver, the LOCR is translated horizontally, using the motion-controlled platform, and the optical power is measured and recorded at discrete locations along a horizontal line that passes through the receiver. The acquired powers are then normalized to the maximum and the 50 % points are located. The center of the receiver is then interpolated from the 50 % points, and the LOCR is translated so that the beam is centered horizontally. This process is repeated to center the beam vertically in the receiver, but instead of translating the LOCR vertically, the height of the optical source breadboard is adjusted. Then the horizontal center is located again, to determine whether the source breadboard shifted during vertical adjustment, and to measure the apparent receiver diameter. The distance between the 50 % points defines LOCR's apparent aperture, which is approximately an ellipse with a diameter of 9.5 mm and a height of 8.5 mm. When the alignment is completed, the window is adjusted to Brewster's angle and the laser's incidence angle relative to the body of the LOCR is adjusted.

To adjust the window to Brewster's angle for a particular wavelength of light, the window mount's gimbals are adjusted so that the intensity of the reflection from the window's first surface is nulled. The null is determined visually, for wavelengths that are visible, using a white target. The intensity of the beam on the target is minimized; if the beam is properly polarized and the window is good, the reflection will be completely nulled. With infrared wavelengths, a single-element germanium detector is used to measure the reflected power. The reflection off the back side of the window is not used, because its intensity can change when vacuum is applied; and when in air, its minimum can occur at an angle different from that of the first surface's, because of the window's nonzero wedge angle. The change in transmittance that occurs when vacuum is applied to the back of the window is accounted for in the uncertainty analysis of window transmittance. The relatively thick window used with the LOCR ensures that the reflections from the front and back of the window are well separated.

The incidence angle of the laser light entering the LOCR is important, because if this angle is not correct, the beam's extent can be limited by apertures other than the LOCR's final aperture. To adjust the incidence angle, first the optical receiver is aligned and the window is adjusted to Brewster's angle, then the location of the beam on the window mount's dust cover is observed. The beam should strike the dust cover on its vertical centerline and approximately 4 mm below its horizontal centerline. This vertical offset is caused by displacement of the beam by the window, as shown in Figure 15. A graphical analysis was performed for the figure using Snell's law [18], using the index of refraction for fused silica at a wavelength of 633 nm [19], and the window at Brewster's angle in air. The index for fused silica is adequate because quartz is crystalline silica. Fused silica's index varies with wavelength, from 1.47 at 400 nm to 1.44 at 2,000 nm, and the change in Brewster's angle for the different wavelengths does not significantly change the offset by more than about 1 mm. An earlier prototype of the NIST window mount is shown in the figure, but the differences from the final version do not affect the offset calculation. However, the offset can change if a window with a different thickness is used.

To force the incident beam to strike the desired vertical location on the window's dust cover, the pitch of the LOCR is adjusted by placing metal shims under the LOCR's bridge electronics enclosure. A pitch angle of approximately 7 mrad is required in the current setup. The pitch is necessary because of imperfections in the welds and vacuum seals used in the long string of vacuum components that connect the window assembly to the cryostat (shown in Figure 7), and imperfections in the polarization angle and window mount. Similarly, the horizontal location of the beam on the dust cover is set by adjusting the yaw of the LOCR. Since the adjustments of pitch and yaw angle will likely change where the beam




Figure 15. The NIST Brewster's angle window mount, showing the vertical offset of the incident laser beam.

strikes the optical receiver, the LOCR's receiver must be realigned after the angles are adjusted, and then the window Brewster's angle and the LOCR's incidence angle readjusted as necessary, in an iterative process. Fortunately, the angle at which the beam is launched from the optical source is very reproducible, so the adjustment of incidence angle has to be performed only once, whenever the LOCR's location on the detector platform is changed.

The resulting incidence angle is double-checked during the window transmittance measurement by observing where the transmitted beam is inside the window mount's vacuum flange; the transmitted beam should be at the center of the flange. Also, if the incidence angle is wrong, the window transmittance at some of the points in the nine-point transmittance matrix may differ significantly from the others, because of partial clipping of the beam by the apertures in the window mount.

One unfortunate result of the necessary adjustments of pitch and yaw is that the center of the cryogenic apertures inside the LOCR, shown in Figure 8, are driven out-of-line with the beam's propagation axis. Apertures other than the final aperture therefore partially limit the extent of the beam, so the final aperture is not necessarily the limiting aperture. One result is that the LOCR's apparent aperture is slightly smaller than its final aperture. Since the pitch adjustment is more severe than the yaw adjustment, the vertical width of the apparent aperture is reduced more than the horizontal width, and the LOCR's apparent aperture is therefore approximately elliptical. The error caused by this misalignment is accounted for in the uncertainty of the relative aperture transmittance correction, since conservative bounding circular apertures are used in the LOCR's relative aperture transmittance calculation; bounding apertures of 8 and 10 mm are used when the LOCR's apparent aperture measures 8.5 x 9.5 mm, as described in Section 4.1.5.

Error in matching the beam's polarization angle to the orientation of the slope of the Brewster's angle window can result in the beam being centered on the window mount's dust cover, but not in the mount's vacuum flange. This difficulty arises because the window mount must be driven significantly off the

LOCR's optical axis when its gimbal mount is used to effect a roll. The window mount does not have a roll-angle adjustment; its roll angle is set when the mount is bolted together. To reduce this problem, the slope of the window, when its gimbal mount's adjustments are in their neutral position, must match the beam's polarization angle (i.e., a line that is normal to the window must lie in the plane of the beam's polarization). Since the rotational angle of the window mount is fixed, it is easiest to adjust the beam's polarization axis by adjusting the roll angle of Polarizer 2. However, the polarizer's roll can be adjusted by only a small amount or the dynamic range of the power stabilizer will be significantly reduced, so the window mount must be bolted together accurately. During assembly of the window mount, a precision level is used to accurately set the window mount's orientation, using the flat areas machined into the top and sides of the mount. The flat on the side of the mount is visible in Figure 12.

Once the LOCR is properly aligned, the device's orientation is locked down with four bolts that connect the LOCR to its platform. Then the LOCR's alignment target is aligned to the beam, after translating the LOCR horizontally a known amount (typically 10 cm), using the motion-controlled detector breadboard. Then if the radiometer's alignment is subsequently lost, the alignment target can be used to rapidly re-align the device.

The alignment of the LOCR is performed carefully, but there is always some uncertainty in the resulting location of the laser beam in the optical receiver. If the receiver's absorptance has some spatial nonuniformity, the uncertainty in location results in an uncertainty in the absorbed power. To assess this uncertainty, measurements of the spatial uniformity of the LOCR's optical absorber are performed after the LOCR is fully aligned. The spatial uniformity is estimated by measuring the power in a stable beam, with the beam located at slightly different locations around the receiver's center. After the influence of random power variation, window transmittance and aperture transmittance are removed, the remaining variation in the measured power is due to nonuniformity in the LOCR's optical receiver's absorptance.

Currently, the LOCR receiver's spatial uniformity is estimated by using a small array of samples taken in a horizontal line through the receiver's nominal center. The array is assumed to be a random sampling of the uniformity around the receiver's center, so is analyzed using standard statistical techniques. The current measurement technique and uncertainty analysis is discussed in detail in Appendix G. The measurement technique will likely be improved in the future, primarily by measuring a matrix of points instead of a linear array. Results obtained to date show that there is significant nonuniformity in the receiver's absorption, which appears to be worse when light with shorter wavelengths is used. The nonuniformity is also likely a function of beam diameter, with larger beams sensing less nonuniformity than smaller beams, because the larger beam integrates the responsivity over a larger area than the smaller beam. Because of the dependence on beam profile and wavelength, the LOCR's nonuniformity is measured each time the beam size or wavelength is changed.

4.2.3 Electrical calibration: The LOCR is electrically calibrated once during each optical calibration, in an environment similar to that encountered during the optical calibration. What is calibrated is the amount of heating power that is applied to the LOCR receiver's electrical heater. In the electrical calibration, the error in the LOCR's electronics is measured at the power levels that are used in the optical calibration. The measured error in the amount of power delivered is corrected for in post-processing.

Currently, software provided by the LOCR's manufacturer is used to perform the electrical calibration. The manufacturer provided two programs, "ACRCAL.EXE" and "TCALVER.EXE". The software's operating instructions are well documented in the manufacturer's Operating Instructions [20], so will not be repeated in detail here. Both programs measure the error in the LOCR's electronics; the main difference is that "ACRCAL.EXE" updates the electronics to remove the error, while "TCALVER.EXE" only measures the error. However, the program "ACRCAL.EXE" is not usually used, because it is very

time-consuming; instead “TCALVER.EXE” is used to measure the error at only the power levels that are actually used in the optical calibration. While using “TCALVER.EXE” adds complexity to the post-processing in the form of an electrical calibration correction factor, the correction factor would be required in any case because the system DMM itself also has measurement error that must be corrected.

Both electrical calibration programs use the system DMM to measure the power applied by the LOCR’s electronics, as described in Section 3.2.3. In addition to the error in the applied power that is measured during the LOCR electrical calibration, there is also error in the DMM’s voltage readings themselves, and there is error in the DMM’s measurement of the LOCR’s standard resistor. These errors are all corrected at the same time, in post-processing when the LOCR electrical calibration correction factor is generated. How the LOCR’s electrical calibration correction factor is used in the optical calibration is discussed in Section 4.3.3, and a detailed description of how the correction is calculated is given in Appendix F. Also described in the appendix is a similar correction that is performed for test detectors that utilize the system’s DMM to measure their outputs.

The LOCR electrical calibration must be performed while the LOCR is fully operational. Therefore the LOCR must be at operating temperature and the window assembly must be in place with the gate valve open. The LOCR and its electronics must also be allowed to reach their steady-state temperature for the given environmental conditions and the environmental conditions should be the same as those present during the optical calibrations. Usually the electrical calibration is performed after the optical calibrations, but the electrical calibration can be performed at any time during the calibration process when the above conditions are present and stable.

Performing the electrical calibration requires connecting the system DMM to the LOCR electronics and connecting the LOCR Computer to the DMM’s GPIB interface. Then the calibration program is executed on the LOCR Computer and the operator follows the on-screen prompts. The operator is required to operate manual switches on the LOCR electronics during the calibration, so the process cannot be automated. These programs will have to be implemented on the main process-control computer before the LOCR computer can be eliminated from the system.

4.2.4 Window transmittance measurement: LOCR window transmittance is measured extremely conservatively. Instead of measuring the transmittance at one location on the window one time, we measure a matrix of locations both before and after the calibrations, to account for uncertainty in the window’s alignment and changes in its transmittance that can occur during the calibration. If the window transmittance changes significantly during the calibration, the uncertainty of the calibration results will be increased correspondingly, so the calibration can either be delivered with an accurate but larger uncertainty, or the calibration can be repeated.

The window assembly must be removed from the LOCR for transmittance measurement. A large gate valve isolates the window assembly from the body of the LOCR; the valve allows the assembly to be removed while there is a vacuum in the LOCR. A small pumping port and valve are located between the gate valve and the window assembly, so that the assembly can be evacuated when reattached to the LOCR, before the gate valve is opened. When the window assembly is placed back on the LOCR after transmittance measurement, it is first attached and then evacuated with the turbomolecular pump, before the gate valve is opened again. The pressure inside the window assembly must be reduced using the turbo pump, so that a inrush of air does not occur when the gate valve is opened. A pressure of 0.005 Pa or less should be attained before the gate valve is opened.

The LOCR must first be aligned to the laser beam and the window Brewster’s angle adjusted, before the window assembly is removed from the LOCR for transmittance measurement. Then the dust cover is

placed over the window mount and rotated to its keyed position. A strip of adhesive paper is attached to the front of the dust cover, the precise location where the laser beam strikes the cover is then marked on the paper. When the assembly is removed from the LOCR for transmittance measurement, the laser is aligned to the location marked on the paper. The transmittance is measured in a nine-point rectangular array around the nominal beam location.

When the window is reattached to the LOCR, the laser beam should strike the same location, but imperfections in the window assembly and rotation of the vacuum flange when it is tightened can change the location by a small amount. The change in location change is minimized by using the same o-ring and clamp, in the same orientation that they were in before the window assembly was removed. To compensate for any remaining change in location, the transmittance is measured both before and after the LOCR is used for calibrations, in an initial and final measurement. For the final measurement, the location of the beam on the dust cover is again marked, so that if the location of the beam moved due to the imperfections in the mount, the nominal center location is more accurate in the final measurement than in the initial measurement. The measurements are combined to obtain the transmittance value used in the calibration, and a conservative uncertainty is derived. How the measurements are combined is described in Appendix H. The paper target documents any shift in location of the beam on the window.

4.2.5 Data acquisition: In the main calibration program that runs on the Process Control Computer, the LOCR is treated as a special direct-reading detector, that is manually operated using the LOCR Computer. Data are transferred from the LOCR Computer to the Process Control Computer using a floppy disk. The LOCR is operated using software called “ACR.EXE”, provided by its manufacturer. The manufacturer’s software will eventually be replaced by software in the main calibration program, allowing fully automated operation of the LOCR, as described in Section 3.3. Operation of the manufacturer’s software is described in the User’s Manual provided by the manufacturer [20], so will not be described in detail here.

The operator must enter several parameters into the “ACR.EXE” program, the most significant of which are the operating point and the control loop parameters. The operating point is defined by the desired temperatures of the heat sink and optical receiver. As previously described, the LOCR is normally operated with heat sink and receiver powers of 2 mW each. The regulation temperatures are chosen so that the desired power is delivered. The absolute heat sink power may be varied slightly to adjust the resulting receiver temperature, but a nominal power of 2 mW is always delivered to the receiver for the purpose of electrical calibration, as described in Sections 3.2.3 and 4.2.3.

There are three preset sets of control-loop parameters that were determined by the manufacturer and are called “fast,” “medium,” and “slow.” The “slow” controller parameters require the longest amount of time for the applied power to reach steady-state, but then the power readings have the lowest noise. With the “fast” parameters, the steady-state power is reached more quickly, but the subsequent power readings are more noisy. The “medium” parameters are a compromise between the “fast” and “slow” parameters. These parameters can be modified by the operator, but to date only the preset values only have been used.

Which control loop parameters are used with the optical receiver’s temperature controller depends primarily on the applied optical power. A large optical power requires a large change in the electrical power applied to the optical receiver, while a relatively small optical power requires little change in the receiver’s electrical power. The period from when a change in the optical power occurs to when the applied electrical power stabilizes is called the settling period. In general, a settling period of 40 s is desired. This settling period provides an adequate balance between how much time the measurement takes and how much the absolute optical power drifts between measurements. When a relatively large optical power is applied, the electrical power that is applied to the optical receiver must change

significantly. For example, to measure 1 mW of optical power, the electrical receiver power must drop from 2.0 to 1.0 mW; a change of 50 %. Similarly, when a relatively small optical power is applied, the electrical power has to change relatively little. For example, to measure 0.1 mW of optical power, the electrical power changes from 2.0 to 1.9 mW; a change of only 5 %.

When the change in electrical power is small, the applied electrical power stabilizes within the desired settling period even when the “slow” controller parameters are used, and the subsequent power readings have the lowest possible noise. However when the change in electrical power is large, the electrical power will take more than the desired settling period to stabilize when the “slow” parameters are used, so the “fast” parameters are used instead, at the expense of increased noise in the measurement of electrical power. The increased noise due to using the “fast” settings with large optical powers is not a problem because while the noise is increased, the signal is also increased, so roughly the same signal-to-noise ratio is obtained with both power levels. The “medium” controller parameters are used with intermediate optical powers as necessary, the parameters used are determined experimentally at the time of the calibration by applying the optical power in a test run and observing the resulting settling period.

The temperature-controller parameters used with the heat sink’s temperature controller must also be selected. Limited experimentation with different heat-sink parameters has been performed with the general conclusion that the system works adequately when the same parameters are used with both controllers. So when the receiver is used with its “fast” parameters, the heat sink is also used with its “fast” parameters. Note that the actual values of the controller parameters, determined by the manufacturer, are different for the two temperature controllers because the heat sink’s thermal response is different from the receiver’s.

One additional parameter must be adjusted by the operator at the time of the optical calibration. This parameter, the phase delay, is adjusted in real time while observing the AC bridge signals with an oscilloscope. The bridge phase is adjusted so that the AC component in the oscilloscope trace is minimized. The phase delay of both the receiver and heat-sink temperature controllers must be adjusted; the adjustment is stable so long as the environmental conditions are stable. Several other parameters are available, such as the AC bridge frequency, but the default values for these parameters work adequately.

When optical power is measured, the ACR.EXE program requires the operator to initiate the measurement, and to open and close the shutter at the appropriate times. The main Process Control Computer also requires operator input at these times; it takes the file names and keeps track of the timing to coordinate its operation with that of the LOCR Computer. After a set of optical calibrations are performed, the operator transfers the acquired data from the LOCR Computer to the Process Control Computer using a floppy disk, and the Process Control Computer converts the data to its internal format. The acquired data are stored in spreadsheet format along with the data for the other detectors and sensors, and preliminary calibration results are calculated.

The LOCR data are analyzed with the dual-baseline technique which is described later, in Section 4.3.1. Basically the technique requires acquisition of the detector’s baseline before the detector is exposed to optical power, acquisition of data while the optical power is applied, and acquisition of a second baseline taken after the optical power is removed. The ACR.EXE program averages the baselines internally and returns only the average value and standard deviation for the baselines, so the actual data acquired during the baselines are unavailable. Data acquired when the LOCR is exposed to the optical power are however available. But since the baseline points are unavailable, the LOCR data are stored in the spreadsheet with a format slightly different from that used for other detectors. The primary difference is that in the spreadsheet the ACR.EXE program’s average for the baselines is used instead of calculating the baselines directly from the raw data.

The period over which data are sampled after stabilizing is called the rating period. With the LOCR, a rating period long enough for the acquisition of 30 samples is usually used. Since the ACR.EXE program samples the LOCR at a rate of approximately once every 1.23 s, the rating period lasts for about 37 s. Therefore, 30 samples are used in the calculation of the average baselines and illuminated power level.

4.2.6 LOCR maintenance: Several of the systems needed to support the LOCR require periodic maintenance. Many of the maintenance steps are necessary to prevent potential damage to the equipment or operator injury. Some maintenance steps are very important, such as inspection of the liquid-helium reservoir vent, which is needed to prevent a potentially catastrophic failure. The required maintenance steps and safety procedures are documented in the High Accuracy Calibration System Operating Instructions [14], so only the most significant events are covered here. However when the LOCR is not in operation, it requires practically no maintenance. Shutting down the LOCR on the weekends therefore eliminates any requirements for weekend maintenance, in addition to providing the other benefits previously mentioned.

As described in Section 4.2.1, periodic monitoring of the cryogenic and related systems is necessary to protect the equipment from potential damage from ice buildup. The liquid-helium and liquid-nitrogen reservoirs' fill ports must be checked for ice formation at least twice a day. If an ice plug forms in the liquid-helium reservoir's fill port and the plug is not removed when the vapor barrier is removed, it is necessary to push the plug into the reservoir using the copper rod previously described. Ice that forms on the outside of the ports is removed by hand, but the heat gun can be used if necessary.

In addition to checking the liquid cryogen fill ports for ice, the reservoirs must also be periodically refilled. The level of liquid nitrogen in the reservoir is checked visually, by looking into one of the liquid-nitrogen fill ports with a flashlight; the liquid nitrogen looks like boiling water. A liquid-nitrogen level sensor is available, but it is less convenient than directly inspecting the level by eye. The liquid helium cannot be seen since it is obscured by a cloud of cold gas, so its level is monitored using the superconducting level sensor that is integrated into the liquid-helium reservoir's cap. The gauge is activated for only short periods because it produces a small amount of heat when active, which consumes liquid helium. The rate of cryogen consumption is very predictable, so the level checks are performed primarily to check that the consumption is as expected. Checking the levels at least twice a day also provides an opportunity to inspect the fill ports for ice.

Refilling the liquid-nitrogen and liquid-helium reservoirs is essentially the same as the initial fill, described in Section 4.2.1. The main difference is that the liquid-helium reservoir pre-cooling step is not needed, since the reservoir is already at liquid-helium temperature. As in the initial fill, the top of the cryostat tends to accumulate ice during the cryogen refill. Therefore the top of the cryostat must be monitored during the refill and any ice buildup removed. The pipe heater attached to the cryostat should remove the ice, but additional heat can be applied by using the heat gun as necessary to keep the cryostat ice-free.

In general, the liquid-nitrogen will be consumed in about 17 hours, so the reservoir is refilled at least twice a day. The liquid helium lasts two to three days; since the LOCR is shut down on the weekend, this rate of consumption usually requires only one liquid-helium refill per week. However high room temperatures, excessive motion, and low levels of liquid nitrogen can cause the liquid helium to be consumed faster; in which case two refills may be needed during the week to continue operation on Friday.

To determine when cryogen storage containers need to be replaced, the amount of liquid remaining in the containers must be monitored. The liquid-nitrogen containers usually have an integrated level gauge, and

the amount of liquid remaining can also be inferred by the container's weight. However, the liquid-helium container does not have a level sensor, so a device called a thistle tube is usually used to measure the liquid level. Like the transfer line, the thistle tube is inserted into the container's fill port. To prevent a buildup of ice inside the port, the level should be measured as quickly as possible. There is little change in the liquid-helium container's weight from full to empty, so its weight cannot be used to judge the liquid level. But an empty container can be quickly identified by checking its primary pressure relief valve; if no gas escapes when the valve is held open, the container is empty.

The vacuum systems used with the LOCR also require periodic maintenance. The turbomolecular pump station is the single most expensive and delicate vacuum component. The turbomolecular pump contains a turbine that turns at a very high-speed on a bearing that is intended to be periodically replaced. The bearing, located at the bottom of the turbomolecular pump's body, contains a pad that is impregnated with a high-quality lubricant. The manufacturer recommends that pad be replaced periodically because the lubricant breaks down with use and time and loses its protective properties, damaging the bearing itself. The bearing in the calibration system's station has not yet been changed, but the pump is used only occasionally and currently has been operated for much less than 2,000 hours. The lubricating oil in the pump station's rotary-vane backing pump must also be replaced periodically. The oil in the pump station has been changed once and the remaining supply of new oil is stored with the pump station. The waste oil is toxic, so must be disposed of properly. Last, there is an oil-trap filter located on the rotary vane pump's exhaust that should also be changed periodically. A spare filter is kept with the station, but the original has not yet been worn enough to warrant replacement. Located on top of the pumping station is a log in which the pump station's usage and maintenance is recorded. Consult the pump station's documentation for details about the replacement intervals.

In addition to the manufacturer's maintenance requirements stated above, care is also taken to maintain the cleanliness of the vacuum components. To prevent the accumulation of oils on the vacuum lines and seals, the vacuum components are always handled with gloves and never with bare hands. Care is also taken to avoid scratching the sealing surfaces. To prevent the buildup of water vapor on the components, after the pump is used, it is vented with dry nitrogen gas, and the vacuum parts are sealed. Keeping the pump clean results in shorter pump-down time and a lower initial pressure. Similarly, whenever the vacuum inside the cryostat is released, the cryostat itself is vented using dry N_2 gas. Care must be taken while venting to prevent pressurizing the cryostat, which could activate its over-pressure burst valve.

To prevent damaging the LOCR's ion pump, it is important that the pump is turned off before the cryogenic liquids in reservoirs expire. When the liquids expire, the reservoir temperature will rise very quickly and the gasses that were cryopumped by the cold reservoirs is released. The resulting sudden increase in pressure can damage the ion pump. The pump's power supply has safety features that shut the pump down before a catastrophic failure, but the electrodes inside the pump will be prematurely worn by the relatively high pressure, so it is important that the pump be manually shut off before the cryogenics expire.

4.3 Calibration Factor Calculation

Two types of optical detector are supported in the calibration system: power and energy meters. Pulse-power and pulse-energy meters are not supported, so all the meters mentioned here are assumed to be continuous-wave (CW) meters, or chopped CW meters. Power meters measure the time-averaged optical power to which the detector is exposed; energy meters measure the total energy injected into the detector. Since the calibration system's absolute power is very stable, the energy can be accurately determined from the average power, by multiplying the average power by the period over which the power was applied. Conversely, the average power can be determined from the total energy by dividing the total

energy by the period over which the power was applied. In the high-accuracy calibration system, the average power is measured. The period over which the power was applied is also accurately measured, for conversion from average power to total energy, and vice versa, as necessary.

4.3.1 Meter output analysis: Three methods of determining the average power from a power meter's output are supported; the single-baseline, dual-baseline, and no-baseline methods. The two methods that use baselines assume that at the start of data acquisition the meter is in steady-state with only background illumination. All three methods use three timing parameters called the sampling period, the settling period, and the rating period. The sampling period is the period between samples of the meter's output. The settling period is the minimum period allowed for the meter to reach steady-state after a change in the applied optical power occurs. The settling period is usually several meter time constants long, so that the meter rises sufficiently close to its ultimate steady-state output. The rating period is the minimum period during which the detector's output is measured after reaching steady state. With some extremely slow detectors, waiting several time constants is not practical, so the customer specifies a shorter settling period. A short settling period is reasonable so long as the meter is consistently used with the shorter interval. The requested settling and rating periods are minima, because the actual number of samples acquired is the requested settling period divided by the sampling period, rounded up to the next whole number. So if a settling period of 1 s is requested and the sampling period is 4 s, a settling period of 4 s is used.

The meter output readings that are acquired during the rating period are averaged and the average is used in the responsivity calculation. The average meter output (O) is given by

$$O = \frac{1}{N} \sum_{i=1}^N o_i, \quad (2)$$

where there are N samples of the meter's output o , within the rating period. Depending on the meter, the average output can have units such as volts, amps, watts, or simply counts. Output readings are recorded during the settling periods, but the values are not used in the power calculations.

The dual-baseline method is used most often, because it compensates for linear drift in the meter's output. With this method, one rating period of samples is acquired before the shutter is opened; the average output during this period is called the initial baseline, or O_{B1} . After the optical power has been applied for a settling period, another rating period of samples is acquired; the average output during this time period is called the signal, O_s . Then the shutter is closed, another settling period passes, then the final rating period of samples is acquired. The average output during this time period is called the final baseline or O_{B2} .

In the dual-baseline method, the baseline during the illuminated period is interpolated from the bracketing baseline measurements; then the change in the meter's output due to the illumination, called O_M , is given by

$$O_M = O_s - \frac{O_{B1} + O_{B2}}{2}. \quad (3)$$

The dual-baseline method is the method usually used with the LOCR, since the two baselines help compensate for its drift with room temperature.

The single baseline method is used with some meters, typically because they have negligible drift, respond slowly and so have a long settling period, have asymmetrical rise and fall times, or do not return to the same baseline after the illumination. The single-baseline method is essentially the same as the dual-baseline method, except that data acquisition is terminated when the shutter is closed, and

$$O_M = O_S - O_{B1}. \quad (4)$$

With the no-baseline method, the shutter is opened at the start of data acquisition. Then one settling period is waited and one rating period of samples is acquired for measurement of the signal; then the shutter is closed and the acquisition is terminated. Therefore no baselines are acquired and the change in the meter's output is simply the signal, so $O_M = O_S$.

Energy-meter output is processed differently from power-meter output. Their outputs are typically a single number that represents the amount of energy measured. Laser calorimeters use a string of equally spaced samples that are processed using the corrected rise algorithm [4]. Performing the corrected rise calculation is complicated, so will not be repeated here; but the data for the corrected rise are acquired in a manner that is essentially identical to the dual-baseline method that is used with power meters. In either case, the resulting energy meter output is a number, called O_M , that is a measure of the applied energy.

The test meter's output O_M is proportional to the average applied power (or energy), but does not necessarily have units of watts (or joules). A scale factor called G may be used to convert the meter output to the desired units. For example, a photodiode-based detector produces an output current that is proportional to the applied optical power. If its output current is amplified using a current-to-voltage converter, the meter's measured O_M will have units of volts. But the customer may want the detector calibrated in terms of amps per watt, so that the calibration factor is independent of the amplifier used. To convert the voltage output back to a current, the measured output voltage is divided by the converter's gain G (in volts per amp, in this case), to give the meter output in terms of amps. Whether or not a scale factor G is necessary is determined by the units of the calibration factor that the customer desires. If units other than the raw units of O_M are desired, an appropriate gain term G must be used. Similarly, if NIST electronics such as a voltmeter are used to measure the meter's output, a correction for error in the NIST electronics, called k_T , is assessed. Since G is usually an amplifier gain, the scaled and corrected detector output, called O_D , is usually given by

$$O_D = k_T O_M / G. \quad (5)$$

If a scale factor is not needed, G is assigned a value of 1, or left out entirely, so that $O_D = k_T O_M$. Similarly, if NIST electrical measurements are not used with the test meter, the term k_T is omitted and $O_D = O_M$.

The absolute power that is measured by the detector is usually desired, so G is usually chosen to give O_D in terms of watts. In the case of the LOCR primary standard, O_M already has units of watts, so $G = 1$ and the measured power, called P , is given by: $P = k_T O_D = k_T O_M$. Special correction factors for the LOCR, described below, are functionally equivalent to the correction factor k_T .

4.3.2 Applied power calculation: The absolute average power that is applied to the test meter is interpolated from bracketing power measurements that are performed with the standard. The interpolation is done in time, using the midpoints of the power injection intervals as the time references. The bracketing primary standard power measurements, called P_1 and P_2 , are acquired at times t_1 and t_2 ,

respectively. The test meter was illuminated at time t_M , where $t_1 < t_M < t_2$. Using linear interpolation, the average power at the time of the test meter illumination, called P_S , is given by

$$P_S = P_1 + (P_2 - P_1) \frac{t_M - t_1}{t_2 - t_1}. \quad (6)$$

The powers P_1 and P_2 are measured at the LOCR's optical absorber. The derived power P_S is therefore the power that would be measured by the LOCR at the time of the test meter calibration. To determine the absolute power that the test meter is exposed to, several correction factors are needed. The LOCR window transmittance (T_W) is the ratio of the optical power transmitted by the LOCR's window, to the amount of power incident on the window. The LOCR receiver absorptance (A_R) is the ratio of the amount of optical power absorbed by the LOCR's optical receiver, to the incident power. The relative aperture transmittance (T_A) is the ratio of the amount of power that is transmitted by the test meter's entrance aperture, to the power transmitted by the LOCR's entrance aperture. The LOCR electrical calibration correction factor (k_L) corrects for the offset in the LOCR's electrical power measurement. Therefore, the power that is applied to the test meter (P_A) is given by

$$P_A = P_S \frac{T_A k_L}{T_W A_R}. \quad (7)$$

The LOCR correction factors T_W , T_A , and k_L are derived, respectively, from measurements of the LOCR's window transmittance, the LOCR's relative aperture transmittance, and DMM and LOCR electronics electrical calibrations. The correction factor A_R is derived from previous measurements and other research, not from measurements performed at the time of the calibration.

The period during which the optical power was injected into the test meter is called the injection period (t_I). The injection period is measured using a precision time-interval counter that is connected to the optical shutter's controller. When an energy meter is calibrated, the applied energy, given by $E_A = P_A \cdot t_I$, is used in place of the average power.

4.3.3 Calibration factor calculation: To determine the calibration factor (CF) for a power meter, usually the detector's output (O_D) is divided by the applied power (P_A), so

$$CF = \frac{O_D}{P_A} = \frac{O_M T_W A_R k_T}{P_S T_A k_L G}. \quad (8)$$

The calibration factor for energy meters is similar, where the applied power (P_A) is replaced by the applied energy (E_A), so

$$CF = \frac{O_D}{E_A} = \frac{O_M T_W A_R k_T}{t_I P_S T_A k_L G}. \quad (9)$$

Equations (8) and (9) give the calibration factor in terms of detector output per unit power (or energy). Therefore, if the meter output is **divided** by CF , the resulting values will agree with those of the NIST measurement system. However, some customers desire the calibration factor in terms of the applied power (or energy) per unit of the detector's output. Then, the detector output is **multiplied** by the calibration factor. In this case, the reciprocal of eq (8) or (9) is used to determine the calibration factor.

The specific equation used is modified as necessary to give the calibration factor in the units that the customer desires. In any case, the equation used is defined explicitly in the calibration report.

5. UNCERTAINTY ANALYSIS

To achieve high accuracy, it is desirable to minimize all the sources of uncertainty as much as possible. But to keep the service affordable to the customer, uncertainty must also be measured in a reasonable amount of time. So in some cases a trade-off is made between the lowest possible uncertainty and an uncertainty that can be estimated with a reasonable amount of time and effort. Towards this end, we may accept a more conservative but simpler uncertainty analysis, even though a smaller uncertainty may be obtainable with more effort. There may be insufficient information to perform a more accurate analysis, or a complex data analysis involved, so a conservative but simpler estimate of uncertainty may be used instead. In some cases, even conservative estimates are so small that they are negligible compared to other terms, so putting more effort into a better analysis would not be advantageous. Additional measurements and analysis may be performed in the future to reduce the dominant uncertainties.

The analysis for many of the uncertainty terms is provided in appendixes. The analysis is located in appendixes to (1) keep the body of this document a reasonable size, and (2) because if the details of the measurements and analysis change in the future, only the appendixes will need to be updated. Only the higher-level concepts are discussed here, with details left to the appendixes.

Uncertainty estimates for the values given in the NIST calibration report are assessed using the following technique, which is derived from the official NIST guidelines [21]. Each component of uncertainty that contributes to the uncertainty of the measurement result is represented by an estimated standard deviation, termed the standard uncertainty u . The combined standard uncertainty u_c , which represents the estimated standard deviation of the result, is obtained by combining the individual standard uncertainties. The standard uncertainties are estimated from Type A evaluations, where magnitudes are obtained statistically from a series of measurements; and from Type B evaluations, whose magnitudes are based on scientific judgement using all the relevant information available.

The Type A uncertainty components are assumed to be independent and, consequently, the standard deviation, S_r , for each component is

$$S_r = \sqrt{\frac{\sum_i x_i^2 - \frac{1}{N} \left(\sum_i x_i \right)^2}{N - 1}}, \quad (10)$$

where the x_i values represent the individual measurements and N is the number of x_i values used for a particular Type A component. The standard uncertainty of each Type A component is the standard deviation of the mean, given by $S_r/N^{1/2}$. The total standard uncertainty of the Type A components is $[\sum (S_r^2/N)]^{1/2}$, where the summation is carried out over all Type A components.

Most of the Type B components are assumed to be independent and have rectangular or uniform distributions (that is, each has an equal probability of being within the region, $\pm\delta_s$, and zero probability of being outside that region). If the distribution is rectangular, the standard uncertainty for each Type B component is equal to $\delta_s/3^{1/2}$, and the total standard uncertainty for all Type B components is $(\sum \delta_s^2)^{1/2}$, where the summation is carried out for all Type B components.

The combined standard uncertainty is determined by combining the total Type A and Type B standard uncertainties in quadrature; the expanded uncertainty is obtained by multiplying this result by a coverage factor of $k = 2$. The expanded uncertainty U is then

Table 1. Example values for the LOCR primary standard's uncertainty components.

Component of Uncertainty	Correction Factor	Standard Uncertainty (%)
Window Transmittance (T_w, u_{TW})	0.99962	0.0138
Receiver Absorptance (A_R, u_{AR})	0.99987	0.0070
Receiver Alignment (u_{LA})	1.00000	0.0075
Electrical Calibration (k_L, u_{kL})	0.99998	0.0034
Heating Inequivalence (u_{HI})	1.00000	0.0002
Combined Standard Uncertainty (u_{PS})		0.0175

$$U = 2\sqrt{\sum_s \sigma_s^2 + \sum_r \frac{S_r^2}{N}}. \quad (11)$$

The values used to calculate the NIST expanded uncertainty (shown in Table 1) are listed in Table 2. The number of decimal places used in reporting the mean values of the calibration factor listed in Table 1 was determined by expressing the expanded NIST uncertainty to two significant digits.

5.1 Uncertainty of the LOCR Primary Standard

Only calibrations based on the LOCR primary standard are considered here; the uncertainty that results from the use of a transfer or other primary standard is not discussed. The uncertainty terms for other standards can be very different from the LOCR's. So if a standard other than the LOCR is used with the system, additional uncertainty analysis and documentation are necessary.

The uncertainties assessed for the LOCR are the result of years of experimentation and analysis and may differ from the uncertainty terms used in other laboratories for cryogenic radiometer-based calibrations. But the LOCR's combined standard uncertainty is consistent with published results from other quality laboratories [22]. When operated at an optical power of 1 mW, The LOCR usually meets its design goal of absolute power measurement with a standard uncertainty of 0.01 %.

The combined standard uncertainty of LOCR power measurements varies from calibration to calibration, because several of the uncertainty components are measured when the calibration is performed. After a large number of calibrations have been performed, a statistical analysis of the measured uncertainties can be performed to determine probable values. Measurement of the uncertainties each time a calibration is performed provides an additional quality assurance check. The combined standard uncertainty of LOCR power measurement has typically varied from 0.01 % to 0.04 %. The combined standard uncertainty cannot be calculated until the calibration runs are complete and after the final window transmittance is measured.

In each LOCR power measurement, the absolute optical power P_L applied to the LOCR is given by the LOCR's raw power measurement multiplied by the LOCR's electrical calibration factor k_L and divided by the product of the LOCR's window transmittance T_w and the LOCR's receiver absorptance A_R . Additional uncorrected uncertainty components that contribute to the uncertainty in the LOCR's optical power measurement are the uncertainty in the LOCR's receiver alignment u_{LA} , and that of the LOCR's

heating inequivalence u_{HI} . Example LOCR-based power measurement uncertainties are given in Table 1 and the uncertainty components are described below. The combined standard uncertainty u_{PL} of the LOCR primary standard's power measurement is given by

$$u_{\text{PL}} = \sqrt{u_{\text{TW}}^2 + u_{\text{AR}}^2 + u_{\text{LA}}^2 + u_{\text{kL}}^2 + u_{\text{HI}}^2}. \quad (12)$$

5.1.1 Window transmittance: The LOCR's window transmittance correction factor is measured. The window transmittance is always less than one because some of the optical power is lost in the LOCR's window assembly, through absorption, scattering, or reflection. The window transmittance is the ratio of the power transmitted to the optical receiver when the window assembly is in the beam, to the power transmitted to the receiver when the assembly is removed.

To measure its transmittance, the window assembly is removed from the LOCR. To measure the power with and without the window in the beam, a test detector is used in place of the LOCR's optical receiver. The measurement is complicated by the fact that whenever the window assembly is reattached to the LOCR after a measurement of transmittance, the vacuum flanges may compress differently, causing the location where the light strikes the window to change slightly. To account for this effect, and for the difficulty of aligning the window during transmittance measurement exactly like it is aligned on the LOCR, a matrix of transmittance measurements is performed. And because the window's transmittance can change during the calibration, its transmittance is measured both before and after the calibrations are performed.

Currently, a nine-point square matrix, centered about the nominal beam location, is acquired in each measurement and a simple statistical analysis is performed. The measurements are combined by use of the technique described in Appendix H. In the future, uncertainty in the window transmittance may be reduced by using a center-weighted technique, described in Appendix I. Also, to reduce the uncertainty in the window's alignment, the visual alignment technique currently used may be replaced with an alignment technique that uses a translatable pinhole aperture detector.

The window transmittance at each location is measured with a technique that combines the window absorbance, scattering loss, and reflectance into a single measurement of window transmittance [3]. The technique relies on duplicating the field of view and reflection characteristics of the LOCR's optical receiver with the test detector, duplicating the window mount alignment, and assuming that the window's properties do not change when it is removed from the cryogenic radiometer's vacuum. The result of the analysis of transmittance is the absolute window transmittance T_{W} and its standard uncertainty u_{TW} , derived in Appendix H.

5.1.2 Receiver absorptance: Receiver absorptance is a correctable imperfection that occurs because some of the power that enters the LOCR's optical receiver is not absorbed and converted to heat. A small amount of the incident optical power is reflected from the receiver. The LOCR's receiver is designed to absorb the vast majority of the collimated laser power, but a small amount is lost. The amount of power that is lost depends primarily on the wavelength of the light; the effect of beam profile and receiver alignment are covered in the relative aperture transmittance correction and receiver alignment uncertainties, respectively.

The LOCR's manufacturer provided an accurate measurement of the receiver's absorptance at a wavelength of 633 nm. They also provided measurements of the reflectance of the black paint used on the receiver, over a wide range of wavelengths. Since the paint's reflectance is essentially constant over the range of wavelengths considered, it is reasonable to assume that the receiver's absorptance is also

constant over the given wavelength range, and many laboratories do consider the absorptance to be constant. But actual measurement of receiver absorptance, performed on a receiver similar to the LOCR's, show that the absorptance is indeed wavelength-dependent [9].

In keeping with our desire to provide conservative uncertainty estimates, the results of the actual absorptance measurements are used to assess the absorptance of the LOCR's optical receiver at wavelengths other than 633 nm. The receiver absorptance correction factor A_R and its associated standard uncertainty u_{AR} are derived in detail in Appendix J. To reduce the uncertainty in the future, the LOCR may be disassembled and its optical receiver's absorptance measured directly. The direct measurement should be performed periodically anyway, because the receiver absorptance may change with time as the paint ages.

Ice can build up on the optical receiver during use because of atmospheric gasses leaking through the LOCR's imperfect vacuum seals and freezing on the cold surface. Ice deposited on the receiver would seriously affect its absorptance. To avoid this problem, the LOCR is warmed to room temperature each week. The LOCR's carbon getter is also purged by the warming.

5.1.3 Receiver alignment: The term receiver (or LOCR) alignment describes the uncertainty that results from the inability to perfectly align the LOCR's optical receiver during use and during measurement of the receiver absorptance. There is insufficient information about the receiver's spatial uniformity and the receiver alignments to make a correction. The uncertainty can also vary with beam diameter and wavelength.

Measuring the LOCR receiver's spatial uniformity is extremely difficult, because of the device's slow response time, the long time it takes to translate the beam vertically, and the system's imperfect stability of beam power and environment. The result is drift in the measurement of relative responsivity that usually exceeds the underlying uniformity being measured. Therefore, rather than attempting to map the receiver's uniformity, its uniformity is estimated by using a small number of samples taken at discrete locations spaced horizontally on the receiver. The technique currently used is described in Appendix G; a value for the uncertainty u_{LA} due to the LOCR's alignment is returned. The measurement is performed each time the LOCR is aligned, so it also serves as a check on quality assurance.

5.1.4 Electrical calibration: The LOCR measures optical power by electrical substitution, so the uncertainty in its electrical measurements directly impact the uncertainty in the measured optical power. The error in the electrical measurements is determined by calibrating the electronics, so that the electrical measurements can be corrected. A single multiplicative correction factor k_L with standard uncertainty u_{kL} is used to correct the optical power reading. Calculation of the correction factor and its uncertainty is described in Appendix F.

The electrical calibration correction factor and its uncertainty are usually calculated in post-processing, after the calibrations are complete. The magnitude of the correction and its uncertainty are usually negligible compared to other components, except for when the LOCR is used with low optical powers. Since the error in the LOCR's electronics is measured for each calibration, it also provides a check on quality assurance for the LOCR because any large change in the error would signify a potential problem with the radiometer itself.

5.1.5 Electrical heating inequivalence: The manufacturer measured the inequivalence between electrical and optical heating by comparing the response when electrical power is delivered to the main heater, which is close to where the optical power is absorbed, to the response when electrical power is delivered to the spare heater at the front of the optical receiver. The electrical power required by each heater to

achieve the same receiver temperature was compared. The manufacturer performed five measurements; the mean disagreement between the two heaters was 0.0017 %, with a standard deviation of 0.002 %.

This manufacturer's measurement is consistent with values used by most other calibration laboratories, but some leave it out entirely [22]. Thermal modeling of the LOCR [11, 12] showed that the actual worst-case electrical heating inequivalence is less than 0.0002 %, which is much smaller than the uncertainty in the inequivalence measurement itself. Therefore the result from the thermal modeling is used as the standard uncertainty of the heating inequivalence, u_{HI} .

5.1.6 Other LOCR uncertainties: Other sources of uncertainty in the LOCR are negligible when compared to the dominant uncertainties, or are random and are therefore included in the measurement repeatability. Two such sources of uncertainty that can be significant, digital quantization and nonlinear drift, are described below.

Because they use a finite number of bits, all digital systems have a finite quantization error, and the LOCR is no exception. If the readings are properly rounded, the quantization error is bounded by plus or minus half of the least significant bit. However, in systems designed like the LOCR, the quantization interval is smaller than the noise in the system, so that the noise randomizes the quantization. A simple but conservative approach to estimating the quantization error is used. The device's illuminated output signal is analyzed by eye and if significant quantization is observed, then a worst-case uncertainty for the quantization error is assessed. But if the variation in the signal is significantly larger than the apparent quantization interval, the quantization is assumed to be randomized and is therefore included in the measurement repeatability, so no additional uncertainty for the quantization is assessed.

A typical illuminated output signal from the LOCR is shown in Figure 16; the variation in the signal is clearly much larger than the apparent quantization interval of 10 nW, so the quantization is sufficiently randomized and no quantization uncertainty is assessed. Quantization uncertainty is also assessed for the test detector, as described in section 5.3.3.




Figure 16. Typical LOCR illuminated output signal. The signal is dithered, so the quantization error is randomized.

The LOCR's background power readings drift with room temperature and other environmental influences, as discussed in section 3.2. Therefore the dual-baseline measurement technique, discussed in section 4.3.1, is used with the LOCR; the technique corrects for linear drift. However if there are components of higher order in the drift, the interpolated baseline can be in error, causing error in the calculated power. If the resulting error in the calculated power is not randomized, it will not be reduced by averaging multiple measurements. For example, if a higher-order drift is always in the same direction, a systematic error can result even with the dual-baseline method. However, if for example the room temperature is drifting randomly, the error in the interpolated baseline will also be randomized, and is therefore included in the measurement repeatability.

To help ensure that the drift is random or at least minimized, the laboratory environment is carefully controlled (when possible) and monitored. The laboratory air temperature, a significant contributor to LOCR drift, is controlled. Uncontrolled parameters, such as the humidity and pressure, are monitored; and if a significant non-random drift is encountered during a calibration, additional calibration runs should be performed. Another source of drift in the LOCR is drift in the temperature of its liquid-nitrogen-cooled heat shield. The liquid nitrogen temperature shield is vapor-cooled, so its temperature is sensitive to the level of the liquid. This drift accelerates rapidly as the liquid-nitrogen expires and in the past has interfered with measurements. So care must be taken to not operate the LOCR when its liquid-nitrogen level is becoming low.

5.2 Uncertainty of the Applied Power

Equation (12) gives the combined standard uncertainty u_{ps} of each individual LOCR power measurement. As described in eq (6), two power measurements are interpolated linearly in time to determine the power that the LOCR would have measured at the time the test detector was illuminated. Ideally, the applied optical power is stable, so $P_A = P_1 = P_2 = P_L$, and the times are irrelevant. To simplify the uncertainty analysis, the power is assumed to be stable and any remaining uncertainty random, so that the uncertainty in the interpolated power is the same as u_{PL} , described in section 5.1. Equation (7) describes how the interpolated LOCR power P_A is used to calculate the power that is applied to the test detector. The uncertainty in the applied power is the uncertainty in the LOCR power measurements combined with the uncertainty of the new terms u_{TA} and u_{RP} , which are described below. The resulting standard uncertainty of the power applied to the test detector u_{AP} is given by

$$u_{AP} = \sqrt{u_{PS}^2 + u_{TA}^2 + u_{RP}^2} = \sqrt{u_{TW}^2 + u_{AR}^2 + u_{LA}^2 + u_{ML}^2 + u_{HI}^2 + u_{TA}^2 + u_{RP}^2}. \quad (13)$$

5.2.1 Relative aperture transmittance: The relative aperture transmittance is a correctable error caused by (a) imperfections in the approximately Gaussian beam that is used in the system and (b) the finite limiting apertures of the detectors. The correction compensates for any difference in the amount of optical power that is applied to the standard and test detectors. Estimation of the correction factor T_A and its standard uncertainty u_{TA} are described in Appendix E.

5.2.2 Measurement repeatability of the applied power: The measurement repeatability is an uncorrectable error, with an uncertainty that is calculated using a Type A evaluation of the results of multiple applied power measurements, which are assumed to vary randomly. Therefore, the mean of the measurements of P_A is the best estimate of the absolute power that was applied to the test detector, and the standard deviation of the mean of the measurements of P_A is the standard uncertainty. The standard uncertainty u_{RP} is the standard deviation of the mean.

5.2.3 Uncertainty of the applied energy: When an energy meter is calibrated, the uncertainty in the applied energy is used. The uncertainty in the applied energy is the same as that for the applied power except that the uncertainty u_{IT} of the measured injection period is also included in the quadrature sum given in eq (13), and the measurement repeatability of the applied energy is used in place of the measurement repeatability of the applied power.

5.3 Uncertainty of the Test Detector Calibration

The uncertainty in the test detector's calibration factor essentially the uncertainty in the applied power, combined with additional uncertainties that are associated with the test detector. However, because the random component of uncertainty in the calibration factor can be smaller than the random component of uncertainty in the applied power, the uncertainty in the calibration factor can be larger or smaller than the uncertainty in the applied power. Equation (8) gives the calculated calibration factor (CF) for an example power meter.

In accordance with our calibration philosophy, uncertainties inherent to the test meter are not evaluated. But there are three additional cases in which an uncertainty will be assessed, because the uncertainty is due to NIST equipment or imperfections in the calibration system. The three cases are for: amplifier gain uncertainty u_G ; uncertainty u_{TE} in the test detector electrical measurement; and uncertainty u_{DQ} due to quantization in the test meter's output. These uncertainties are not appropriate for all test meters; gain uncertainty is an issue only when a NIST amplifier is used, electrical measurement uncertainty is appropriate only if NIST equipment is used to measure the test detector's output, and quantization is an issue only with digital test meters.

When the component of uncertainty is inappropriate for the particular test meter, its uncertainty is assigned a value of zero, or the component is left out completely. The combined standard uncertainty u_{CF} of the calibration factor is

$$\begin{aligned} u_{CF} &= \sqrt{u_{PS}^2 + u_{TA}^2 + u_G^2 + u_{TE}^2 + u_{DQ}^2 + u_{RC}^2} \\ &= \sqrt{u_{TW}^2 + u_{AR}^2 + u_{LA}^2 + u_{ML}^2 + u_{HI}^2 + u_{TA}^2 + u_G^2 + u_{TE}^2 + u_{DQ}^2 + u_{RC}^2}. \end{aligned} \quad (14)$$

The new components of uncertainty introduced in eq (14) are described below. The combined standard uncertainty in the applied power and calibration factor for an example test meter calibration are shown in Table 2. The example is from an actual 100 μ W calibration of a test detector with a digital output, so gain and electrical measurement uncertainty are inappropriate, but a small uncertainty due to quantization is assessed.

5.3.1 NIST amplifier gain: When a NIST amplifier is used to amplify or convert a test detector's output, the uncertainty u_G of the amplifier's gain is assessed. For example, a NIST current-to-voltage converter is frequently used to amplify and convert the output current from a photodiode-based detector to a voltage. Since the converter is not an actual part of the customer's meter and is therefore not always used with the detector, its uncertainty must be assessed.

Parameters describing the NIST amplifier are also included in the calibration report. For example, the current-to-voltage converter can load the detector so its input characteristics, such as input impedance and voltage burden, are quoted. These parameters are measured by NIST and an upper bound is stated. For

Table 2. Example uncertainties for the applied power and calibration factor calculation.

Component of Uncertainty	Correction Factor	Standard Uncertainty (%)
Window Transmittance (T_w, u_{TW})	0.99962	0.0138
Receiver Absorptance (A_R, u_{AR})	0.99987	0.0070
LOC Receiver Alignment (u_{LA})	1.00000	0.0075
Electrical Calibration (k_L, u_{kL})	0.99998	0.0034
Heating Inequivalence (u_{HI})	1.00000	0.0002
Relative Aperture Transmittance (T_A, u_{TA})	1.00000	0.0011
Measurement Repeatability of the Applied Power (u_{RP})*		0.0072
Combined Standard Uncertainty of the Applied Power (u_{AP})		0.0190
Ammeter Quantization (u_{DQ})	1.00000	0.0007
Measurement Repeatability of the Calibration Factor (u_{RC})*		0.0075
Combined Standard Uncertainty of the Calibration Factor (u_{CF})		0.0191

* - Type A evaluation with $N = 6$.

the example NIST current-to-voltage converter, when the converter is used with gains of up to 100,000 V/A, the input impedance is less than 0.1 Ω , and the voltage burden is less than 1 mV.

5.3.2 NIST electrical measurement: If the test detector produces an analog electrical signal such as a voltage that is measured by NIST electrical equipment, the uncertainty in the NIST equipment must be assessed. For example, consider a test detector that produces an analog voltage that is measured by the calibration system's DMM. Evaluation of the uncertainty in the system's DMM is described in relation to the LOC electrical calibration in Appendix F; assessment of test detector electrical measurement uncertainty is similar. The error and the uncertainty in the voltage measurements are used to obtain the correction factor k_T and uncertainty u_{kT} for the measured detector output.

For example, consider a test detector that produces an analog voltage that is processed with the dual-baseline method. To simplify the analysis, the two baselines are assumed to be equal; in practice the two baselines are usually very close together. The baselines have an average voltage v_B and the illuminated signal has an average voltage v_S . The voltmeter has an absolute error e_B with uncertainty u_B when measuring voltage v_B , and an absolute error e_S with uncertainty of u_S when measuring voltage v_S . These parameters are determined from the voltmeter's calibration, as described in Appendix F. The voltmeter error is such that the actual voltage is the measured voltage **plus** the voltage error.

When the dual-baseline method is used, the meter output is calculated using eq (3). Given the voltmeter errors, a multiplicative correction factor can be calculated:

$$k_T = 1 + \frac{e_S - e_B}{v_S - v_B}. \quad (15)$$

The uncertainties u_S and u_B are assumed to be uncorrelated, which is reasonable if $|v_S| \gg |v_B|$. Even though two baselines are used, the baselines are ideally identical, so the averaging of the two baselines in eq (3) is ignored and the uncertainty in the baseline is simply u_B . Therefore, the combined standard

uncertainty u_{KT} in the correction factor k_T is given by

$$u_{KT} = \frac{\sqrt{u_S^2 + u_B^2}}{v_S - v_B}. \quad (16)$$

When using the single-baseline method from eq (4), the correction factor and uncertainty are the same as with the dual-baseline method, because the baseline averaging in the dual-baseline method is ignored. When using the no-baseline method, only the voltage measurement error at v_S is relevant, so the correction factor becomes

$$k_T = 1 + e_S / v_S. \quad (17)$$

When estimating the uncertainty for the no-baseline method, a term for the voltage offset in the measurement system's cabling and connectors must be assessed. The baseline methods are differential measurements, so the offset is canceled; but with the no-baseline method (and any other nondifferential voltage measurements), the offset must be assessed. To ensure the quality of the system's cabling and connectors, the offset is measured periodically, as a quality-control measure. In the measurements, when a high-quality electrical short terminates the cables, the absolute value of the offset voltage is consistently below 1 μV . Therefore, using a worst-case bound of $\pm 10 \mu\text{V}$ for the offset voltage is probably very conservative, yet does not significantly overestimate the uncertainty because the absolute value of v_S is usually greater than 1 V. The voltage offset is not corrected, because it is not measured before each calibration and can change with time and temperature. The resulting standard uncertainty in k_T is

$$u_{KT} = \frac{1}{v_S} \sqrt{u_S^2 + 3.33 \cdot 10^{-11} \text{V}^2}. \quad (18)$$

5.3.3 Quantization in digital test meters: Because of the calibration system's power stability, quantization in digital test meters can lead to uncertainty in its power measurement if there is insufficient dithering in the meter itself. As described in section 5.1.6, the LOCR's output is sufficiently dithered, so that any error due to quantization is randomized. The test meter's output must also be inspected, to determine whether an uncertainty for its quantization must be assessed.

To determine whether quantization is an issue, the meter's sampled output is observed. In general, if a meter outputs a digital power reading with a fixed number of digits, it may have quantization error. But meters with an analog output can also be quantized, so a more general rule is used: if the meter's output is quantized, the sampled readings will occur at discrete levels, with no readings occurring between the levels. Determining whether or not the output is sufficiently dithered is more speculative, but in accordance with the desire to return a conservative estimate of uncertainty, it is better to include the uncertainty when unnecessary than to omit it when necessary. So a simple but conservative rule is used to determine whether quantization error should be assessed: if the samples in the rating period are all at one or two levels, an uncertainty for the potential quantization error should be assessed. The worst-case scenario is when all the sampled readings are at a single level, with a standard deviation of zero.

Figure 17 shows the sampled readings from a test meter with potential quantization error. Since all the readings occur at only two levels, the output is considered insufficiently dithered to randomize the quantization error, so an uncertainty for the quantization is assessed. In the example, the power level was intentionally adjusted so that two levels are visible. Because about half of the samples occur at each level in the example, the uncertainty in the average is actually less than the assessed quantization uncertainty. However, had the applied optical power been slightly different, all of the readings could have occurred at




Figure 17. The output current from an illuminated photodiode, showing two discrete quantization levels.

a single level, then the assessed uncertainty due to quantization is accurate. To be conservative, an uncertainty for the quantization is assessed, even though in cases like that shown in Figure 17, the actual uncertainty due to quantization is reduced by averaging.

A bounded uncertainty estimate is used to quantify the quantization error. The meter's output is assumed to be properly rounded to the nearest quantization interval. Therefore, the largest possible error in the quantized output reading is one-half of the quantization interval Q . The quantization error has a rectangular distribution, so the standard uncertainty due to quantization, u_Q , is

$$u_Q = \frac{Q}{2\sqrt{3}}. \quad (19)$$

Depending on the output processing method used, the samples from multiple rating periods may have to be observed. In all methods, the rating period samples acquired when the detector is illuminated must be considered. When a processing method with baselines is used, quantization in the baseline rating periods must also be considered.

When the no-baseline processing method is used, the standard uncertainty due to quantization is equal to that for the illuminated rating period only, which is given by eq (19). When the single-baseline method is used, the uncertainty for both the illuminated and background rating periods must be considered. Using eq (19), define u_{QB} as the standard uncertainty due to quantization in the baseline rating periods and u_{QS} as the uncertainty in the illuminated rating period. Both uncertainties are uncorrelated and since the baseline is subtracted from the signal, the combined standard uncertainty, u_{DQ} , is the quadrature sum of the standard uncertainties u_{QB} and u_{QS} . When the dual-baseline method is used, the two baselines are essentially the same, so the standard uncertainty due to quantization is the same as that for the single-baseline method.

The effect of quantization error in the baselines can be reduced by allowing the meter to auto-range; then the size of the quantization interval in the baseline will be much smaller than the quantization interval while illuminated, if the meter has higher ranges of gain available. Then the contribution of the

quantization in the baseline will be insignificant compared to the quantization when the detector is illuminated, so the combined uncertainty is essentially the same as that for the no-baseline method.

5.3.4 Measurement repeatability of the calibration factor: The measurement repeatability of the calibration factor is essentially the same as the measurement repeatability of the applied power, described in section 5.2.2, except that the calculated calibration factors are used in the analysis, instead of the calculated applied powers. The measurement repeatability uncertainty is the result of a Type A evaluation of the measured calibration factors, so the average is the best estimate, and the standard uncertainty is the standard deviation of the mean.

5.3.5 Uncertainty in the calibration factor for energy meters: When an energy meter is calibrated, the uncertainty in the injection period must be included in the uncertainty analysis. The uncertainty in the calibration factor for energy meters is the same as that for power meters, given in eq (14), except that the standard uncertainty in the injection period, called u_{IT} , must also be included in the quadrature sum.

5.4 Other Uncertainties

Many additional parameters are specified in the calibration report to fully characterize the conditions present during the calibration. Each parameter is specified by an estimation with expanded uncertainty, a conservative upper bound, or an approximation based on a model. Parameters that vary during the calibration, such as room temperature, are specified with an uncertainty that includes the variation. The intent of these uncertainty statements is to document the conditions present at the time of the calibration; they only incidentally affect the calibrated quantities, so do not warrant a complex measurement procedure or detailed uncertainty analysis. However, this information must be documented in order for the high-accuracy calibration to be reproduced.

5.4.1 Test detector alignment: The uncertainty in the centering of the laser beam in the test detector's entrance aperture is specified, along with a measure of the beam's incidence angle in relation to the test detector's entrance aperture. The orientation of the detector relative to the beam's polarization angle is also specified. A specific alignment procedure, described in section 4.1.4, is used so that the alignment can be performed reproducibly.

The uncertainty in the centering of the laser beam in the test detector's entrance aperture is based on experiments performed by the laboratory operator and can differ for other operators. In the experiment, the operator aligns a detector a few times and the variation in the beam's final location on the detector is determined. A conservative upper bound for the error in the alignment is then derived, by selecting a value significantly larger than any variation that occurred in the experiment. The resulting upper bound is given in the calibration report. A more formal experiment would have to be performed and documented to obtain accurate standard uncertainties for the alignment.

Three different centering techniques are used, each with its own uncertainty bound. The coarsest technique involves visually interpolating the center of the entrance aperture using an external scale such as a ruler. Experiments have shown this type of alignment is performed with an uncertainty of less than 1 mm, for a beam of 1 to 2 mm diameter (in general, smaller beams can be centered more accurately than larger beams). If the detector has a visual alignment aid such as a crosshair on its aperture cover, the beam can be centered more accurately. Experiments have shown a centering accuracy of 0.5 mm is conservative for beams of 1 to 2 mm diameter when such alignment aids are used. The uncertainty of the cross hair's centering on the aperture cover, and slop in the aperture cover itself, must also be considered, but are usually less than the operator's visual limitations.

A better alignment technique uses a small aperture as an alignment aid, such as a small hole in the center of the aperture cover, or the hole through the center of an attached optical fiber adapter. The beam is centered over the hole by using the detector itself to peak the power transmitted through the hole. Again, the accuracy of the hole's centering in the aperture cover must also be included in the uncertainty assessment. Experiments have shown when using the common screw-on FC optical fiber adapter as the alignment aid, the detector is centered to better than 0.2 mm when a beam of 1 to 2 mm diameter is used. Most of this uncertainty with this technique is due to uncertainty in the centering of the FC connector in the screw-on adapter and slop in its screw threads. Therefore it is best to use the same FC adapter in the calibration that is used with the detector in practice.

The result of the centering uncertainty analysis is given as a bound in a statement such as: "With power meter x, the beam was aligned visually to the center of the aperture cover's cross hairs, with an uncertainty of less than 0.5 mm." The description is changed to match the actual alignment technique used, as appropriate. If the customer specifies a particular alignment procedure, their procedure is used and no uncertainties are assessed.

The incidence angle of the laser beam as it enters the test detector is also reported. Usually, the incidence angle is set as close to normal incidence as possible. But to sufficiently document the conditions during the calibration, the difference in the beam's incidence angle from the normal relative to a surface on the detector is measured and stated in the report. The surface on the detector that is used to judge what is normal, such as the detector's face plate or entrance aperture, is also specified.

When a detector that is intended for use in optical fiber systems is aligned, the beam is centered using the hole in its optical fiber adapter, and the incidence angle is adjusted using the normal for the fiber adapter itself. Then the calibration system's collimated beam travels straight down the center of the fiber adapter and strikes the detector's active surface at the same location that a fiber-coupled beam would.

The error in the incidence angle is determined by holding a glass flat against the specified detector's surface and locating the back reflection from the glass flat in relation to the incident beam. Then a simple geometric calculation is performed to determine the difference in the incidence angle from the detector surface's normal. Conservative uncertainty bounds for the geometric distance measurements are used to determine bounds for the error in the incidence angle calculation, from which the standard uncertainty is derived using a Type B evaluation, and the expanded uncertainty with a coverage factor of 2 is specified. The result is given in a statement such as: "the beam's incidence angle was approximately 0.7 ± 0.2 mrad from normal relative to the detector's entrance aperture." Information about the direction of the error in the incidence angle is not currently provided. But in general, the error in the incidence angle is usually in the vertical direction, because no pitch adjustment is available with most detector mounts. If an exceptionally accurate incidence angle is required, such as when an optical fiber adapter is used as previously described, a special mount with pitch adjustment can be used with the detector.

The laser beam polarization angle's relation to the test detector is described by specifying the detector's orientation relative to horizontal. The polarization angle is described as vertical. Some characteristic of the detector is used to describe its orientation on the table, such as: "the detector was mounted so that the label on its faceplate was up" or "the detector was mounted on a post holder, with its threaded mounting screw down." The relation between the polarization angle and the test detector's orientation is thus defined, but no numerical measure is given. Currently the actual polarization angle is not known accurately, so the precise polarization angle relative to the detector orientation cannot be specified.

5.4.2 Laser beam spatial characteristics: To fully characterize the laser beam used in the calibration, several spatial beam parameters are specified in the calibration report. A theoretical Gaussian beam

model is combined with measured parameters to provide values for beam's diameter and divergence at the detector's entrance aperture, and theoretical values for the Gaussian beam waist's diameter and location.

Since some of these parameters are theoretical values that are not directly measured, they are described as “approximate” in the calibration report, and no uncertainty is specified. The uncertainty of the theoretical values is dependent on how well the actual laser beam matches a true Gaussian beam, which is quantified by the beam's M^2 parameter [23]. Currently, the beam's M^2 parameter is not routinely measured, but occasional measurement of the theoretical waist diameter and far-field divergence have been consistent with the theoretical values to better than 10 %. So using a standard uncertainty of 10 % with the parameters should be conservative, but the figure cannot be guaranteed without measurement, so approximate values are reported with no estimate of uncertainty.

Using a few measured parameters, the Gaussian beam model for the system's laser beam is derived. The Gaussian beam model is described in Appendix D. The model returns the final beam waist diameter and location relative to the detector plane, and the beam's characteristics in the detector plane.

The beam's diameter is measured using an autonomous scanning-slit beam profiler, which measures the beam's $1/e^2$ intensity diameter. The device is relatively noisy, so the uncertainty in measurements of its beam diameter is dominated by random error, which is evaluated using a Type A analysis. Therefore the mean of the diameter measurements is used as the measured beam diameter, called D , and the standard uncertainty that results from the random variation in the measurement, called u_A , is the standard deviation of the mean.

Additional error in the scanning-slit measurement technique is caused by the ratio of the slit's width to the beam diameter [10]. The error in the scanning-slit diameter measurement is less than 0.17 % when:

$$w < 0.05D. \quad (20)$$

Where w is the slit width and D is the $1/e^2$ intensity beam diameter. The error of 0.17 % is one-sided, the width measured with the scanning slit is always larger than the actual width. And the error is smaller when the ratio of the slit width to beam diameter is smaller, so the error of 0.17 % is the upper bound on the error, so long as eq (20) is true.

Highly accurate measurements of beam diameter are not necessary, so the error in the diameter measurement that results from the finite slit width is not corrected; instead the maximum error is used as an uncertainty bound. Therefore, when eq (20) holds true, the standard uncertainty u_s due to the finite slit diameter is

$$u_s = 0.0017 / \sqrt{3} \approx 0.001. \quad (21)$$

Additionally, the manufacturer provides the uncertainty u_M of the scanning-slit device's calibration. For example, since the beam profiler's slit width is 5 μm , beams with a diameter of only 0.1 mm can be measured without increasing u_s . The uncertainties u_A , u_s , and u_M are uncorrelated, so the combined standard uncertainty in D , called u_D , is given by

$$u_D = \sqrt{u_A^2 + (Du_s)^2 + u_M^2}. \quad (22)$$

Because of imperfections in the polarized laser beam and coma distortion caused by the tilt in the spatial filter lenses, the beam produced by the optical source is usually elliptical, with its long axis in the horizontal direction. Both vertical and horizontal diameters are measured, and the average of the vertical and horizontal diameters is the nominal beam diameter.

Call the two diameter measurements D_V and D_H , and their associated standard uncertainties u_{DV} and u_{DH} , each given by eq (22). Then the average diameter is $D_A = (D_V + D_H) / 2$. We assume that the uncertainties u_{DV} and u_{DH} are uncorrelated, and that D_V and D_H are the bounds of a rectangular distribution. The standard uncertainty u_{DA} in D_A is given by

$$u_{DA} = \frac{1}{2} \sqrt{u_{DV}^2 + u_{DH}^2 + \left(\frac{D_V - D_H}{\sqrt{3}} \right)^2} \quad (23)$$

5.4.3 Laser beam wavelength and spectrum: Uncertainty in the wavelength measurement is dependent on the device that is used to measure the wavelength. The precision wavelength meter has the lowest uncertainty, but can be used only with sources that produce a relatively narrow spectral line. The wavelength of broader sources and sources that produce multiple spectral lines is measured using an optical spectral analyzer.

The precision wavelength meter automatically determines its inherent accuracy, and infers the optical source's linewidth and mode structure. The device limits the resolution of its wavelength display in accordance with its inherent accuracy. Using the displayed resolution as an uncertainty bound, the equivalent standard uncertainty u_l inherent in the devices wavelength measurement is the displayed resolution divided by $\sqrt{3}$.

The wavelength measurement repeatability uncertainty u_R is also assessed using a Type A evaluation of multiple wavelength measurements. If the source's center wavelength is unstable, or the source has multiple spectral lines, u_R can be prohibitively large, so the optical spectrum analyzer is used to measure the wavelength instead. Since the uncertainty components are uncorrelated, the combined standard uncertainty u_λ for the wavelength measurement is the quadrature sum of u_l and u_R .

The wavelength meter's display resolution, standard uncertainty, and implied optical source bandwidth are listed in Table 3. The wavelength meter measures and corrects for the air's index and can return the wavelength in air or vacuum; the wavelength in vacuum is reported. Where Δv is the implied optical

Table 3. Precision wavelength meter's estimated uncertainty and source bandwidth estimate.

Displayed Resolution (nm)	Estimated Standard Uncertainty u_l (pm)	Implied Optical Source Bandwidth (THz)
0.0001	0.06	≤ 0.0015
0.001	0.58	≤ 0.015
0.01	5.8	≤ 0.15
0.1	58.0	≤ 1.5
1.0	580.0	≤ 15.0

source bandwidth, λ is the vacuum wavelength, and C is the speed of light in vacuum, an upper-bound for the one-half power spectral linewidth in vacuum ($\Delta\lambda$) is given by

$$\Delta\lambda = \lambda \frac{2C}{(2C + \lambda\Delta\nu) - (2C - \lambda\Delta\nu)}. \quad (24)$$

By observing the wavelength meter's "trigger" and "window" output, whether the implied bandwidth is for a single spectral line or multiple spectral lines can be inferred. If there are multiple spectral lines, the implied bandwidth can be for either the dominant spectral line, or for the half-power bandwidth of all the lines combined, depending on their relative intensities. In the former case, a stable wavelength will be returned by the device, but in the latter case, the measured wavelength will probably be erratic, as the mode structure of the multi-spectral-line source changes with time. For such multiline sources, the measurement repeatability of the wavelength meter's measurements will be excessive, so the optical spectrum analyzer is used instead of the wavelength meter. Therefore, the wavelength meter is used for wavelength and bandwidth measurement only when the source has a single spectral line, or multiple lines with one dominant line. The source spectrum is then specified in the report with a statement such as: "The laser radiation was contained in a single spectral line having a half-power linewidth of ($\Delta\lambda \cdot 10^9$) nm or less." Note that a "single spectral line" can actually contain multiple longitudinal modes that are too close together for the wavelength meter to resolve. So the source cannot be stated as having a single longitudinal mode unless it is known to have only a single mode for other reasons, such as the laser's geometry or its feedback structure.

When an optical spectrum analyzer (OSA) is used instead of the wavelength meter, because the source has multiple spectral lines or an unstable center wavelength, the wavelength and linewidth are much more uncertain. The near-infrared wavelength OSA currently used has a standard uncertainty in its wavelength measurements of 0.13 nm. The visible wavelength OSA has a standard uncertainty in its wavelength measurements of 0.29 nm. To determine the combined standard uncertainty for the OSA wavelength measurement, this uncertainty is added in quadrature with the measurement repeatability uncertainty, from a Type A evaluation of multiple measurements. Measurement of spectral linewidth using the OSA are limited by the device's resolution of about 0.1 nm. When used with sources that contain multiple spectral lines, the OSA allows documentation of the structure of the spectral lines, including their relative intensity. The spectral content of such sources is described using the OSA measurements. For example, although a particular 1319 nm laser produces three spectral lines, one is dominant; so an accurate wavelength meter measurement is possible, but the spectral line structure is defined using the OSA analysis. The laser is then described with a statement such as: "The wavelength of the laser is 1319.52 ± 0.13 nm. Within this wavelength region the laser has three spectral lines, each with half-power linewidth of less than 0.1 nm." For lasers with more complex spectra, the verbal description can be supplemented with a graph of the spectrum as measured by the OSA.

5.4.4 Environment specification: The air temperature, pressure, and humidity during the test detector calibration are monitored using the autonomous weather station and specified in the calibration report along with their expanded uncertainties. Also, the index of refraction of the air surrounding the test detector is also calculated from these parameters and specified in the report.

The autonomous weather station records readings once every minute. The weather station measures two air temperatures, the temperature at the weather station's location, and the temperature in the detector plane where its external temperature sensor is placed. Using its internal sensors, temperature, humidity and pressure are measured near the optical source, on the opposite side of the scattered light shield from the detectors. Pressure and humidity do not vary significantly between the two locations, but the

temperature can. The temperature difference between the detector plane and the optical source is used to estimate the size of the temperature gradient u_G in the vicinity of the detectors. The uncertainty in the weather station's measurements is determined by calibrating the instrument, as described in Appendix K. The variation in the parameter measured during the calibration is included in the uncertainty, as described in the appendix.

The numerical refractive index of the air is necessary for calculating the air wavelength of the laser light, and can affect the test detector's absorptance. Because of the altitude of the laboratory, the air's index of refraction is typically 1.00022, which is significantly lower than usually assumed. The index of the air around the detectors is calculated from the detector plane air temperature, pressure, relative humidity, and the laser's wavelength [5]. Other parameters used in the calculation, enhancement factor (1.00364), saturation water vapor pressure (2693 Pa), and compressibility (0.999686), are based on typical local values for pressure (85,000 Pa), temperature (22° C), relative humidity (30 %), and CO₂ abundance (0.0003). The resulting index is relatively insensitive to these parameters, so using the typical values is sufficient.

The uncertainty in the index calculation is determined using the law of propagation of uncertainties [21], assuming that the components of uncertainty are uncorrelated. The equations used are complex, so numerical partial derivatives are used instead of an analytic solution. Example values and the result of the numerical partial derivative calculation are given in Table 4. The nominal index and its expanded uncertainty are given in the calibration report. Using a coverage factor of $k = 2$, the resulting index of refraction for the example parameters is 1.0002239 ± 0.0000023 . Uncertainty in the model used in the index calculation is much smaller, typically around $2 \cdot 10^{-8}$ [5], so can be ignored; the uncertainty in the measured temperature, pressure, and humidity dominate.

Table 4. Components of uncertainty used in the calculation of the index of refraction of the air.

Uncertainty Component	Nominal Value	Partial Derivative	Standard Uncertainty
Temperature (K)	298.11	$-7.5 \cdot 10^{-7}$	0.34
Pressure (Pa)	83879.0	$2.7 \cdot 10^{-9}$	406.0
Relative Humidity (%)	30.5	$-1.1 \cdot 10^{-8}$	5.0

REFERENCES

Note. Brand and company names are used here only so that the system can be duplicated, and imply neither endorsement by the National Institute of Standards and Technology (NIST), nor that the materials or products identified are necessarily the best available for the purpose.

- [1] J. E. Martin, N. P. Fox, and P. J. Key, "A cryogenic radiometer for absolute radiometric measurements," *Metrologia* **21**, 147-155 (1985).
- [2] T. R. Gentile, J. M. Houston, J. E. Hardis, C. L. Cromer, and A. C. Parr, "The NIST high-accuracy cryogenic radiometer," *Appl. Opt.* **35**, 1056-1068 (1996).
- [3] P. V. Foukal, C. Hoyt, H. Kochling, and P. Miller, "Cryogenic absolute radiometers as laboratory irradiance standards, remote sensing detectors, and pyroheliometers," *Appl. Opt.* **29**, 988-993 (1990).
- [4] E. D. West, "Data analysis for isoperibol laser calorimetry," Nat. Bur. Stand. (U.S.) Technical Note 396 (1971).
- [5] F. E. Jones, "The refractivity of air," Nat. Bur. Stand. (U.S.) J. Res. **86**, 27-32 (1981).
- [6] D. J. Livigni, C. L. Cromer, T. R. Scott, B. C. Johnson, and Z. M. Zhang, "Thermal characterization of a cryogenic radiometer and comparison with a laser calorimeter," *Metrologia* **35**, 819-827 (1998).
- [7] E. Hecht, *Optics*, 3rd ed., pp. 604-610 (Addison Wesley Longman, New York, NY, 1998).
- [8] M. Young, "Mode-field diameter of single-mode fiber by far-field scanning," *Appl. Opt.* **37**, 5605-5619 (1998).
- [9] G. E. Chamberlain, P. A. Simpson, and R. L. Smith, "Improvements in a calorimeter for high-power CW lasers," *IEEE Trans. instrum. meas.* **IM-27**, 81-86 (1978).
- [10] D. J. Livigni and X. Li, "Spatial uniformity of optical detector responsivity," *Proc. NCSL Workshop and Symposium*, Chicago, IL., 337-352 (1994).
- [11] B. C. Johnson, *Heat transfer analysis and modeling of a cryogenic radiometer*, M.S. Thes. University of Florida (1997).
- [12] B. C. Johnson, A. R. Kumar, Z. M. Zhang, D. J. Livigni, C. L. Cromer, and T. R. Scott, "Heat transfer analysis and modeling of a cryogenic laser radiometer," *J. Thermophys. Heat Trans.* **12**, 575-581 (1998).
- [13] Cambridge Research and Instrumentation, Inc., 21 Erie Street, Cambridge, MA 02139, USA.
- [14] D. J. Livigni, "High-Accuracy Calibration System Operating Instructions," unpublished, Natl. Inst. Stand. Technol. (1998).
- [15] L. D. Dickson, "Ronchi ruling method for measuring Gaussian beam diameter," *Opt. Eng.* **18**, 70-75 (1979).

- [16] J. M. Houston, and D. J. Livigni, "Comparison of two cryogenic radiometers at NIST," Natl. Inst. Stand. Technol. J. Res. **106**(4): 641 (2001).
- [17] O. R. Touayar, J. M. Coutin, and J. Bastie, "Measurement of the reflectance of the INM cryogenic radiometer cavity at several wavelengths," Metrologia **35**, 387-391 (1998).
- [18] M. Young, *Optics and Lasers including Fibers and Optical Waveguides*, 4th ed., Chapter 2, Eq. (2.1) (Springer, New York, 1994).
- [19] Malitson, I.H. "Interspecimen comparison of the refractive index of fused silica," J. Opt. Soc. Am. **55**(10), 1205–1209 (October 1965).
- [20] "Cryogenic Radiometer Operating Instructions," Cambridge Research and Instrumentation, Inc., 21 Erie Street, Cambridge, MA 02139, USA.
- [21] B. N. Taylor and C. E. Kuyatt, "Guidelines for evaluating and expressing the uncertainty in NIST measurement results," Natl. Inst. Stand. Technol. Technical Note **1297** (1994).
- [22] R. Goebel, M. Stock, and R. Köhler, "Report on the international comparison of cryogenic radiometers, based on transfer detectors," BIPM Report BIPM-2000/9 (2000).
- [23] "Lasers and laser-related equipment -- Test methods for laser beam parameters -- Beam widths, divergence angle and beam propagation factor," ISO 11146 (1999).
- [24] Ref. 7, pp. 249-250.
- [25] Ref. 18, Section 2.2.4, pp. 14-15.
- [26] J. T. Verdeyen, *Laser Electronics*, 3rd ed., Section 3.6. Verydin, ABCD matrix, pp. 76-79, (Prentice Hall, NJ, 1995).
- [27] F. Barnes, "ECEN 4045 Lecture Notes," University of Colorado, 1991.
- [28] A. E. Siegman, *Lasers*, pp. 665-670 (University Science Books, CA, 1986).
- [29] Ref. 18, Section 10.2.6, pp. 260-261
- [30] Ref. 7, pp. 587.
- [31] Ref. 18, Section 7.2.4, pp. 158-159.
- [32] Ref. 18, Section 8.3.3, pp. 205-208.
- [33] Ref. 7, Section 4.6, pp. 109-120.
- [34] HP 3458A multimeter operating, programming, and configuration manual, Appendix A, Hewlett-Packard Company (1994).
- [35] Davis Instruments Precision Weather Instruments Catalog, Davis Instruments, Weather Monitor II specifications (1996).

APPENDIX A. Glossary

Some of the more unusual or potentially misunderstood terms and abbreviations used throughout this document are defined. In this section, terms that are defined here are shown in *italics*.

Apparent Aperture: The *apparent aperture* is an estimate of the size and shape of a *detector's* equivalent limiting aperture. The *apparent aperture* is determined using a *spatial uniformity scan* of the *detector*. When the *spatial uniformity scan* data is scaled to 100 % in the *detector's* center, the 50 % contour line describes the *detector's apparent aperture*. With a *detector* that has a true limiting aperture, the *apparent aperture* is consistent with the limiting aperture; at the edge of the aperture, half the beam is blocked and the other half enters the detector, so the response is reduced to 50 % of the response in the detector's center. The apparent aperture accurately describes the limiting aperture when the diameter of the scanning beam is much smaller than the diameter of the limiting aperture; if the beam is too large, the apparent aperture will be somewhat smaller than the true limiting aperture.

However, some *detectors* do not have a true limiting aperture, their entrance aperture is larger than the *detector's* photosensitive element. For example, the apparent aperture of a *detector* with no limiting aperture is shown in Figure 6. The 50 % contour line shown in the figure describes the *detector's apparent aperture*. The *detector* has a deceptively large physical entrance aperture, which is circular, and has a diameter of slightly more than 12 mm. The *apparent aperture* is much smaller and is not circular. Knowledge of the *detector's apparent aperture* is necessary for calculation of the *relative aperture transmittance*.

Detector and Meter: The terms *detector* and *meter* are sometimes used interchangeably to describe a device that measures optical power or energy. Specifically, *detector* refers to the part of the device into which the optical power enters and is sometimes called the “head.” *Meter* specifically refers to the electronic package into which the *detector* plugs, and which provides the power or energy readout. Some devices are a single package so there is no distinction. We make this distinction because we will sometimes calibrate a customer's *detector* using our own *meter* electronics; in this case we calibrate only the *detector* part of the device.

Detector Alignment: When a detector is aligned to the laser beam, the terms pitch, yaw, roll, and translation are used. The terms are depicted in Figure 14. The motions are made relative to the beam's propagation axis, and the direction of down as defined by the direction of gravity, where the beam's propagation axis is perpendicular to the direction of down. A vertical translation is a movement of the detector, without any angular change, along an axis that is parallel to the direction of down. A horizontal translation is a movement, without any angular change, in a direction that is perpendicular to the direction of down. During *detector alignment*, translations are usually performed within the *detector plane*, which is a plane to which the beam propagation axis is normal, at a prescribed location along the propagation axis. A horizontal translation in the detector plane is defined as along the x-axis, with x increasing to the right, when facing the direction of the beam's propagation. A vertical translation in the detector plane is defined as along the y-axis, with y increasing in the direction opposite of down. A translation without rotation that is along the beam's propagation axis is a longitudinal translation, and is considered a translation along the z-axis, with z increasing with the distance from the optical source.

Pitch, yaw, and roll are angular adjustments, in which the detector is rotated about an axis. Pitch is a rotation around an axis that is perpendicular to the direction of the beam's propagation, in the horizontal direction, or perpendicular to the direction of down. Yaw is a rotation around an axis that is parallel to the direction of down. Roll is a rotation around an axis that is parallel to the beam's propagation axis. The location of the rotational axes with respect to the detector is unimportant, as is any translation which

results from the rotations, because when a detector is aligned, the rotations are adjusted before the translations. However to aid in the visualization of the rotations, the rotational axes can be considered to pass through the center of the detector's mass, as shown in Figure 14.

Detector Breadboard: The *detector breadboard* is the large platform to which the LOCR and *test detectors* are mounted. The *detector breadboard* is mounted to the motion-controlled detector stage, which provides translation along the horizontal axis, perpendicular to the propagation axis of the beam.

Detector Plane: The *detector plane* is a plane perpendicular to the propagation axis of the beam, where the limiting or input aperture of the *detectors* are placed. The diameter, profile, and ideally the divergence angle of the laser beam in the detector plane are specified. When the divergence angle is not measured, it is estimated using Gaussian beam analysis. The location of the *detector plane* along the laser beam's propagation axis is defined by the location of the entrance or limiting aperture of the *standard*. A *detector's* entrance aperture is by convention the hole in the *detector's* face-plate where the light enters. With the LOCR *standard*, the best location for the detector plane is not obvious. Currently the *detector plane* is placed where the LOCR's final internal aperture is, because it is the primary limiting aperture.

Device Under Test (DUT): The *device under test* is a generic term that refers to the device being calibrated, whether it is a simple *detector* or complete *meter* of power or energy.

Gaussian Beam: The high-accuracy calibration system uses *Gaussian beams*, produced by single transverse mode (TEM_{00}) lasers or a single-mode optical fiber cores. Ideal *Gaussian beams* have a Gaussian intensity profile, and their propagation through space can be described by specifying the beam's wavelength, direction, waist location, waist diameter, and the refractive index of the surrounding media. The actual beams used in the system are only approximately Gaussian, but *Gaussian Beam* theory effectively predicts the beam's behavior. Measurements performed to confirm the theoretical results show that in the high-accuracy calibration system, the theory predicts the beam diameter to within a few percent. The diameter of a *Gaussian beam* is usually specified by the $1/e^2$ intensity diameter. The quality of a *Gaussian beam* is usually specified with the m^2 parameter, which is equal to unity for a perfect *Gaussian beam* and is greater than unity for imperfect beams.

Laser Power Controller (LPC): The *Laser Power Controller* is a commercial device used to stabilize the laser beam's absolute power. The *LPC* consists of three main parts and their connecting cables; the *LPC* electronics, the *LPC* modulator and the *LPC* monitor. The modulator uses a liquid-crystal-based device to rotate the polarization axis of the transmitted laser beam. The beam is then demodulated using a polarizer and the demodulated power is measured with the monitor. The electronics package controls the modulator and monitor, and returns the power measured by the monitor and the modulator's transmission ratio. The *LPC* modulator itself contains an internal polarizer and monitor detector and can function alone without the *LPC* monitor. The external *LPC* monitor is used to compensate for loss and polarization mixing in the spatial filter and other intervening optical elements. The monitor that is used by the *LPC* electronics is determined by the device into which the monitor cable is plugged. Two *LPC's* are used in the high accuracy calibration system, one is designed for use with visible wavelengths, the other is optimized for infrared wavelengths.

Optical Source Breadboard: The large breadboard that supports the optical source optics is called the *optical source breadboard*. The breadboard is supported on five manually operated leveling posts; the leveling posts allow adjustment of the laser beam's height above the laboratory table.

Primary Standard: A *standard* that is directly traceable to physical standards, usually electrical standards. The *primary standard* should be well documented and understood from a physics standpoint and ideally

uses simple processes that can be accurately modeled. Cryogenic radiometers like the LOCR are currently the best primary standard available for optical power measurement, in terms of absolute accuracy.

Relative Aperture Transmittance: The *relative aperture transmittance* for a *test detector* is the ratio of the amount of power transmitted by the *test detector*'s aperture to the amount of power transmitted by the *standard*'s aperture, when the detectors are illuminated by the beam used in the calibration system. The ratio is used to correct the amount of power delivered to the *test detector*, thereby accounting for any differences in the aperture shapes or sizes. The relative aperture transmittance is either measured directly for each *test detector*, or is estimated using measurements performed with circular apertures. The *detector*'s *apparent aperture* is used to estimate its *relative aperture transmittance*.

Sensors: *Sensors* are devices that are sampled between sets of optical detector measurements. They are used for special purposes, such as monitoring a *detector*'s temperature, or measuring a *detector*'s bias voltage. *Sensors* are not required for calibrations.

Spatial Filter: The *spatial filter* is an optical device that consists of two lenses and a pinhole and is used to limit the spatial frequencies present in a beam. The device is commonly used to convert a uniform illumination into an approximately *Gaussian beam*. In the high-accuracy calibration system, the optical source itself produces an approximately *Gaussian beam*, and a *spatial filter* is used to remove scattered light and other artifacts from around the *Gaussian beam*. To produce the desired beam size in the detector plane, the location of the second lens in the *spatial filter* is adjusted slightly around the optimal collimation distance. The result of the lens adjustment is accurately predicted by *Gaussian beam* theory. To further control the beam size and parameters, lenses of different focal lengths can be used.

Spatial Uniformity Scan: The *spatial uniformity scan* of a *detector* is a map of the *detector*'s responsivity to a given beam, versus the location of the beam on the *detector*. The beam location has two dimensions and is in the plane of the *detector*'s entrance aperture. The responsivity is sampled at discrete locations, with a sample spacing fine enough to prevent aliasing of the underlying features. The resulting *spatial uniformity scan* data are usually presented with a wire-frame or contour map; a contour map showing the *apparent aperture* of a *detector* is presented in Figure 6. The data can also be presented as a ray-traced surface or matrix of numbers, and statistical analysis can be performed on the data to quantify the *detector*'s uniformity. The *spatial uniformity scan* data are accurate only when the *detector* is used with a similar beam; the *detector*'s uniformity can be different for other beams.

The spatial uniformity of the *standard* is important in the high-accuracy calibration system, because the *detector*'s alignment is imperfect. Uncertainty in the *detector*'s alignment causes uncertainty in the *detector*'s responsivity; the uncertainty in responsivity that results from the spatial uncertainty is dependent on the *detector*'s spatial uniformity. Since the *detectors* are aligned only once during a calibration, this uncertainty does not contribute randomly, so must be accounted for separately. The uncertainty that results from imperfect alignment of the *test detectors* is not accounted for in the calibration, but should be added by the end-user when the calibrated *test detector* is used as a *standard* in their laboratory.

Standard: The *standard* is the *meter* which is used as the reference in a calibration. The *standard* can be a *primary standard* or a *transfer standard*.

Test Detector: The device that is being calibrated is called the *test detector*, the device used as the reference in the calibrations is the *standard*. Each *test detector* consists ideally of a *meter* and a *detector*,

but the *detector* alone is sometimes calibrated. The *test detector* can be an optical meter of either power or energy.

Transfer Standard: A laboratory *standard* that may not be suitable as a *primary standard*, but is traceable to a *primary standard* through calibration. *Transfer standards* are sometimes used in place of the *primary standard* in calibrations, because they are usually quicker and easier to use, but they also add additional uncertainty to the calibration. *Transfer standards* are sometimes called secondary standards.

Turbomolecular pump: The *turbomolecular pump*, sometimes abbreviated as turbo pump, is a self-contained high-vacuum pumping station that is used as a roughing pump for the LOCR. The station includes (a) a turbomolecular pump, which is a high-speed turbine and is the main pumping mechanism, (b) the control electronics, and (c) a rotary-vane backing pump. Pressure gauges are attached to the vacuum system plumbing above the turbomolecular pump; the plumbing serves as a foreline when the LOCR is evacuated, so the gauges indirectly measure the pressure in the LOCR's cryostat during evacuation. Protection circuitry and electrically controlled valves make the pump safe to use, but care must be taken to prevent damage to the pump and to keep the inside of the vacuum system components clean and dry. Some oil vapor is released by the rotary-vane backing pump, so the pump exhaust is vented to the outside of the building.

APPENDIX B. Instructions for Submitting Meters for Calibration

1. Select the calibration service from the NIST Calibration Group Calibration Services Users Guide SP 250 Appendix Fee Schedule. The guide is available in the following ways:

A. From calibration services office:

Calibration Services
National Institute of Standards and Technology
100 Bureau Drive, Stop 2330
Gaithersburg, MD 20899-2330
301-975-2002 phone
301-869-3548 fax
calibrations@nist.gov

B. From the NIST calibration service web site:

<http://ts.nist.gov/ts/htdocs/230/233/calibrations/>.

The high-accuracy laser power and energy calibration service is Service ID Number 42162S, “Special Test for High Accuracy Laser and Optical Fiber Power Measurements,” which is listed in the “Lasers and Optoelectronic Components Used with Lasers” portion of the “Optical Radiation Measurements” section. The service is listed in chapter 7 section E of the guide, and at the web site:

<http://ts.nist.gov/ts/htdocs/230/233/calibrations/optical-rad/laseroptoelectronic.htm>.

2. Set a calibration date by contacting the personnel responsible for the desired calibration service, given with the description of the calibration services.

3. Specific technical information is required from the customer. If some of the information is omitted, a best-guess assumption will be used and documented in the calibration report. When the high-accuracy calibration service is desired, the following information is requested:

A. The desired laser wavelength. We currently support vacuum wavelengths of 458.06, 476.62, 488.13, 496.65, 514.67, 632.99, 1064.42, 1343.09 and 1550.43 nm.

B. The desired power level in the range from 0.1 to 1.0 mW. Other power levels are negotiable.

C. The desired beam size to be used with the detector. We use approximately circular Gaussian-profile beams. Please specify the $1/e^2$ intensity width desired at the detector’s input aperture. The divergence angle or radius of curvature at the detector’s entrance aperture can also be specified, but the beam size may have to be changed to support the desired divergence.

D. Detector cleanliness is important for high-accuracy calibrations, since a speck of dust on the detector can cause a calibration error larger than the stated uncertainty. So NIST will normally blow out the detector with filtered compressed air. The customer should

inform NIST if their detector can not be cleaned in this way, then the cleaning step will be skipped for their detector.

E. The desired detector orientation. The system uses a polarized beam, so detector orientation can be significant.

F. The desired detector alignment technique. How the detector should be aligned relative to the laser beam should be specified.

G. A complete description of how the meter should be setup before the calibration, such as the desired initialization routine for any coupled electronics.

H. A complete description of how the meter's output data should be read and processed, including desired time intervals. We can use single-, dual-, or no-baseline analysis techniques or the corrected-rise algorithm. The customer should specify which technique to use, and the desired settling, rating, sampling, and injection intervals.

I. Instructions on the use of any ancillary components needed for the meter, such as external temperature controllers or optical choppers.

J. Exceptions to the above, including using other power levels and wavelengths, and using NIST meters and equipment with customer's detector, will be considered on a per-customer basis.

4. Shipping instructions. Meters sent to NIST for calibration should be shipped in well-padded foam or otherwise mechanical-shock-insulated case, appropriate for return shipment back to the customer.

Operation instructions or instruction manuals should be included, as well as the customer-chosen set-up parameters for instrument functions. The customer should include all cables and connectors that are necessary to calibrate the equipment as specified. Exceptions and special cases will be considered on a per-customer basis.

APPENDIX C. Sample Calibration Report

A sample calibration report is shown on the following pages. The report details a calibration where the customer provided three photodiode-based detectors and one ammeter, which were calibrated as a unit (each detector was in turn calibrated using the single ammeter). The report for other types of meters will differ slightly. For confidentiality, places where customer-specific information would be given are shown as a description within square brackets; the symbol # is used to signify where an alpha-numeric entry would occur. The bottom-centered page number is an artifact of this document and would not appear in an actual report.

The calibration equation and uncertainty analysis terms given in the report can change slightly depending on the specifics of the calibration. However, the calibration equation is always consistent with eq (8) or eq (9), depending on whether the meter measures power or energy. In the example, each detector is combined with the ammeter to form a single unit that measures optical power. The ammeter is read through its GPIB output, so no analog voltage is measured. Therefore the calibration equation given in the report is the same as eq (8), with $G = 1$, $O_M = I_M$, and $k_T = 1$. Had a NIST device, such as a current-to-voltage converter with analog voltage output, been used with the detectors, G and k_T would be included in the calibration equation and in the uncertainty analysis.

Other information that changes in each report includes the measured optical power and detector calibration factors, the environmental conditions, the descriptions of the optical source, and the specifics of the detector alignment. Depending on the specifics of the calibration, additional measurements may be performed, such as a bias voltage measurement.

Each report is for calibrations performed at a single wavelength. Calibrations performed at additional wavelengths are reported separately, since each wavelength is a separate calibration, performed at a different time, possibly with a different laser. The environment parameters, beam parameters, detector parameters, and voltmeter uncertainty can therefore differ for each wavelength. To avoid the complexity of detailing all this different information in a single report, calibrations at additional wavelengths are reported separately.



Report of Calibration

for

LASER POWER METER

[Manufacturer name] Model: [model #], Serial No: 1, NIST ID: [#####]

[Manufacturer name] Model: [model #], Serial No: 2, NIST ID: [#####]

[Manufacturer name] Model: [model #], Serial No: 3, NIST ID: [#####]

with [Manufacturer name] PICOAMMETER

Model [#####], Serial No: [#####]

Submitted By:

[Customer Name]

[Customer Address line 1]

[Customer Address line 2]

Measurement Summary

The laser power measurement system supplied by the customer consisted of a picoammeter, a connecting cable with switch box, and three detectors called [model #]. Each detector was calibrated by comparison to the NIST Laser Optimized Cryogenic Radiometer (LOCR). The measurement results are provided in Table 1. The comparison was performed by direct substitution of the standard with the test detector, using the calibration system shown in Figure 1. Each detector was calibrated, in turn,

Table 1: Calibration Results (Wavelength = 633 nm).

Detector Serial Number	<i>N</i>	Applied Power			Result	
		Average (mW)	Expanded Uncertainty (%), (<i>k</i> = 2)	Average (A/W)	Standard Deviation (%)	Expanded Uncertainty (%), (<i>k</i> = 2)
1	12	1.00063	0.023	-0.50937	0.0087	0.023
1	6	0.099671	0.025	-0.50945	0.0127	0.025
2	6	1.00077	0.022	-0.50924	0.0042	0.023
2	6	0.099692	0.023	-0.50932	0.0070	0.024
3	6	1.00085	0.022	-0.50831	0.0071	0.023
3	6	0.099689	0.024	-0.50832	0.0090	0.024

Folder No. & NIST ID: [#####-##] & [#####]

Date of Report: [Report completion date]

Test Number: [#####]

Reference: [Purchase order number, submission date]

Page # of #

LASER POWER METER

Model: [Manufacturer, Model #], Serial No: 1, NIST ID: [#####]

Model: [Manufacturer, Model #], Serial No: 1, NIST ID: [#####]

Model: [Manufacturer, Model #], Serial No: 1, NIST ID: [#####]

[Manufacturer name] Picoammeter, Model [###], Serial No: [#####]

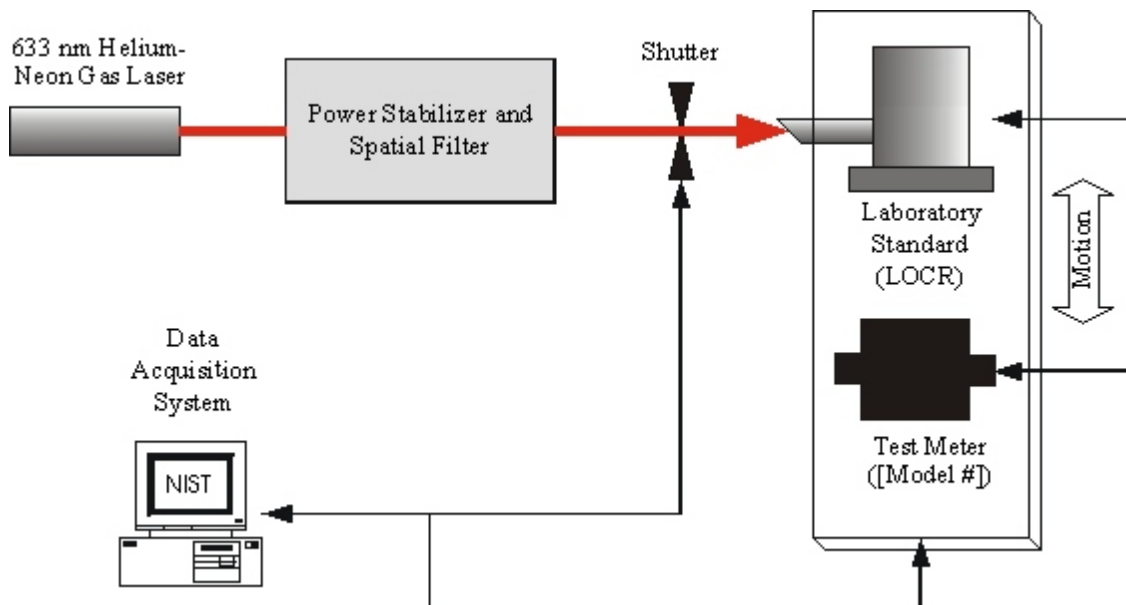


Figure 1. Measurement system diagram.

using the supplied picoammeter and connecting cable with switch box. The calibration system used a power stabilizer to stabilize the absolute power and a spatial filter to remove scattered light. The optical power applied to the test detector was calculated by interpolating between power measurements performed with the standard before and after the test detector measurement, and then applying the appropriate correction factors.

The calibration was performed using a helium-neon gas laser, with a nominal center vacuum wavelength of 632.99217 ± 0.00017 nm. The laser radiation was contained in a single spectral line having a half-power linewidth of less than 0.004 nm. The laser beam was vertically polarized and had an approximately Gaussian intensity profile with $1/e^2$ diameter of 1.97 ± 0.10 mm at the detector's entrance aperture. The full-width beam divergence was approximately 0.7 mrad at the entrance. The beam was approximately equivalent to an ideal Gaussian beam, having a waist of diameter 1.0 mm, located 2.1 m in front of the detector's entrance aperture.

The detectors were oriented so that the NIST ID tag was at the top, and the electrical output connector was on the right side, when facing the detector's entrance aperture. The end of the detector displaying the manufacturer's label was used as the entrance. The top of the detector was leveled to within $\pm 0.2^\circ$. The detectors were aligned to the beam visually by adjusting the detector so that the beam passed through both the front and back irises, when these openings were minimized. The detectors were then finely aligned so that with the irises open, the transmitted beam had a diamond shape with an approximately uniform illumination, and so that the irises closed around the center of the transmitted

Folder No. & NIST ID: [#####-##] & [#####]

Date of Report: [Report completion date]

Test Number: [#####]

Reference: [Purchase order number, submission date]

Page # of #

LASER POWER METER

Model: [Manufacturer, Model #], Serial No: 1, NIST ID: [#####]

Model: [Manufacturer, Model #], Serial No: 1, NIST ID: [#####]

Model: [Manufacturer, Model #], Serial No: 1, NIST ID: [#####]

[Manufacturer name] Picoammeter, Model [###], Serial No: [#####]

beam. All of the switches on the connecting cable's switch box were turned on, except for the bias switch, which was turned off. Both the front and back irises were fully open during the calibrations.

Each detector's output was measured using the [Manufacturer, model #] picoammeter's GPIB port. Three strings were used to initialize the instrument: "L0X", "C0G1T5W0Y3X" and "R0X". These commands activated auto-ranging, among other things. To initiate a reading, the command "X" was sent. Each detector's output was processed with a dual-baseline method using 20 s minimum rating periods and 1 s minimum settling periods. A sampling rate of 1 sample per second was used. The environment near the detectors during the calibration was: temperature = $21.63 \pm 0.72^\circ \text{C}$, pressure = $836 \pm 11 \text{ hPa}$, relative humidity = $35.0 \pm 6.2 \%$, and air index of refraction = 1.0002200 ± 0.0000036 . The tolerances given in the text are expanded uncertainties with a coverage factor of $k = 2$.

The calibration factor for the detector was calculated by dividing the detector's output current by the optical power delivered to the detector. Four correction factors were used: the LOCR window transmittance (T_w), the LOCR receiver absorptance (A_R), the relative aperture transmittance (T_A), and the LOCR electrical calibration (k_L). The detector's calibration factor (CF) in amps/watt is given by the equation:

$$CF = \frac{I_M T_w A_R}{P_s T_A k_L},$$

where P_s is the applied power in watts, interpolated from bracketing primary standard measurements, and I_M is the detector output in amps. If the output current of the detector is **divided** by the calibration factor listed in Table 1, the resulting values will agree with those of the NIST measurement system.

Uncertainty Assessment

Uncertainty estimates for the values given in the NIST calibration report are assessed using the *NIST Guidelines for Evaluating and Expressing the Uncertainty of NIST Measurement Results* (Technical Note 1297). Uncertainty estimates for the values given in the NIST calibration report are assessed using the following technique, which is derived from the official NIST guidelines. Each component of uncertainty that contributes to the uncertainty of the measurement result is represented by an estimated standard deviation, termed the standard uncertainty u . The combined standard uncertainty u_c , which represents the estimated standard deviation of the result, is obtained by combining the individual standard uncertainties. The standard uncertainties are estimated from Type A evaluations, where magnitudes are obtained statistically from a series of measurements; and from Type B evaluations, whose magnitudes are based on scientific judgement using all the relevant information available.

The Type A uncertainty components are assumed to be independent; and consequently, the standard deviation, S_p , for each component is

Folder No. & NIST ID: [#####-##] & [#####]

Date of Report: [Report completion date]

Test Number: [#####]

Reference: [Purchase order number, submission date]

Page # of #

LASER POWER METER

Model: [Manufacturer, Model #], Serial No: 1, NIST ID: [#####]

Model: [Manufacturer, Model #], Serial No: 1, NIST ID: [#####]

Model: [Manufacturer, Model #], Serial No: 1, NIST ID: [#####]

[Manufacturer name] Picoammeter, Model [###], Serial No: [#####]

$$S_r = \sqrt{\frac{\sum_i x_i^2 - \frac{1}{N} \left(\sum_i x_i \right)^2}{N-1}}, \quad (10)$$

where the x_i values represent the individual measurements and N is the number of x_i values used for a particular Type A component. The standard uncertainty of each Type A component is the standard deviation of the mean, given by $S_r/N^{1/2}$. The total standard uncertainty of the Type A components is $[\Sigma(S_r^2/N)]^{1/2}$, where the summation is carried out for all Type A components.

Most of the Type B components are assumed to be independent and have rectangular or uniform distributions (that is, each has an equal probability of being within the region, $\pm\delta_s$, and zero probability of being outside that region). If the distribution is rectangular, the standard uncertainty for each Type B component is equal to $\delta_s/3^{1/2}$, and the total standard uncertainty for all Type B components is $(\Sigma\sigma_s^2)^{1/2}$, where the summation is carried out for all Type B components.

The combined standard uncertainty is determined by combining the Type A and Type B standard deviations in quadrature; the expanded uncertainty is obtained by multiplying this result by a coverage factor of $k = 2$. The expanded uncertainty U is then

$$U = 2 \sqrt{\sum_s \sigma_s^2 + \sum_r \frac{S_r^2}{N}}. \quad (11)$$

The values used to calculate the NIST expanded uncertainty (shown in Table 1) are listed in Table 2. The number of decimal places used in reporting the mean values of the calibration factor listed in Table 1 were determined by expressing the expanded NIST uncertainty to two significant digits.

The correction factors, their normalized standard uncertainties, and the combined standard uncertainty for the applied power and the calibration factor calculation are given in Tables 2 and 3, for nominal power levels of 1.0 and 0.1 mW, respectively. The combined standard uncertainty of the applied power is the quadrature sum of the common uncertainties and the applied power measurement's repeatability; the combined standard uncertainty of the calibration factor is the quadrature sum of the common uncertainties, the calibration factor measurement's repeatability, and the test detector ammeter quantization uncertainty. The two measurement repeatabilities differ slightly, depending on the amount of test detector noise and the effectiveness of the power interpolation. An uncertainty for quantization in the ammeter was assessed, because the ammeter's five-digit readings varied insufficiently to randomize the uncertainty.

Folder No. & NIST ID: [#####-##] & [#####]

Date of Report: [Report completion date]

Test Number: [#####]

Reference: [Purchase order number, submission date]

Page # of #

LASER POWER METER**Model:** [Manufacturer, Model #], Serial No: 1, NIST ID: [#####]**Model:** [Manufacturer, Model #], Serial No: 1, NIST ID: [#####]**Model:** [Manufacturer, Model #], Serial No: 1, NIST ID: [#####]**[Manufacturer name] Picoammeter, Model [###], Serial No: [#####]****Table 2. Uncertainty components for 1 mW calibrations.**

Component of Uncertainty	Detector Serial Number					
	1		2		3	
	Correction Factor	Std. Unc. (%)	Correction Factor	Std. Unc. (%)	Correction Factor	Std. Unc. (%)
LOCR Receiver Absorptance (A_R)	0.99992	0.0002	0.99992	0.0002	0.99992	0.0002
LOCR Window Transmittance (T_W)	0.99970	0.0072	0.99970	0.0072	0.99970	0.0072
LOCR Electrical Calibration (k_L)	1.00000	0.0030	1.00000	0.0030	1.00000	0.0030
LOCR Alignment	1.00000	0.0075	1.00000	0.0075	1.00000	0.0075
LOCR Electrical Heating Inequivalence	1.00000	0.0002	1.00000	0.0002	1.00000	0.0002
Relative Aperture Transmittance (T_{RA})	1.00003	0.0033	1.00003	0.0033	1.00003	0.0033
Measurement Repeatability of the Applied Power		0.0016*		0.0013**		0.0016**
Combined Standard Uncertainty of the Applied Power		0.0115		0.0114		0.0115
Ammeter Quantization	1.00000	0.0006	1.00000	0.0006	1.00000	0.0006
Measurement Repeatability of the Calibration Factor		0.0025*		0.0017**		0.0029**
Combined Standard Uncertainty of the Calibration Factor		0.0116		0.0115		0.0117

* - Type A evaluation with $N = 12$. ** - Type A evaluation with $N = 6$.**Folder No. & NIST ID:** [#####-##] & [#####]**Date of Report:** [Report completion date]**Test Number:** [#####]**Reference:** [Purchase order number, submission date]

Page # of #

LASER POWER METER**Model:** [Manufacturer, Model #], Serial No: 1, NIST ID: [#####]**Model:** [Manufacturer, Model #], Serial No: 1, NIST ID: [#####]**Model:** [Manufacturer, Model #], Serial No: 1, NIST ID: [#####]**[Manufacturer name] Picoammeter, Model [###], Serial No: [#####]****Table 3. Uncertainty components for 0.1 mW calibrations.**

Component of Uncertainty	Detector Serial Number					
	1		2		3	
	Correction Factor	Std. Unc. (%)	Correction Factor	Std. Unc. (%)	Correction Factor	Std. Unc. (%)
LOCR Receiver Absorptance (A_R)	0.99992	0.0002	0.99992	0.0002	0.99992	0.0002
LOCR Window Transmittance (T_W)	0.99970	0.0072	0.99970	0.0072	0.99970	0.0072
LOCR Electrical Calibration (k_L)	1.00000	0.0032	1.00000	0.0032	1.00000	0.0032
LOCR Alignment	1.00000	0.0075	1.00000	0.0075	1.00000	0.0075
LOCR Electrical Heating Inequivalence	1.00000	0.0002	1.00000	0.0002	1.00000	0.0002
Relative Aperture Transmittance (T_{RA})	1.00003	0.0033	1.00003	0.0033	1.00003	0.0033
Measurement Repeatability of the Applied Power		0.0051*		0.0019*		0.0040*
Combined Standard Uncertainty of the Applied Power		0.0125		0.0116		0.0121
Ammeter Quantization		0.0006		0.0006		0.0006
Measurement Repeatability of the Calibration Factor		0.0052*		0.0029*		0.0037*
Combined Standard Uncertainty of the Calibration Factor		0.0125		0.0118		0.0120

* - Type A evaluation with $N = 6$.**Folder No. & NIST ID:** [#####-##] & [#####]**Date of Report:** [Report completion date]**Test Number:** [#####]**Reference:** [Purchase order number, submission date]

Page # of #

LASER POWER METER

Model: [Manufacturer, Model #], Serial No: 1, NIST ID: [#####]

Model: [Manufacturer, Model #], Serial No: 1, NIST ID: [#####]

Model: [Manufacturer, Model #], Serial No: 1, NIST ID: [#####]

[Manufacturer name] Picoammeter, Model [###], Serial No: [#####]

For the Director,
National Institute of Standards and Technology

Report Reviewed By:

Marla L. Dowell, Group Leader
Sources and Detectors Group
Optoelectronics Division

John H. Lehman, Project Leader
Sources and Detectors Group
Optoelectronics Division

Report Prepared / Calibrated By:

David J. Livigni, Electronics Engineer
Sources and Detectors Group
Optoelectronics Division

Folder No. & NIST ID: [#####-##] & [#####]

Date of Report: [Report completion date]

Test Number: [#####]

Reference: [Purchase order number, submission date]

Page # of #

APPENDIX D. Gaussian Beam Model and Spatial Filter Pinhole Selection

The optical elements considered in the Gaussian beam model are shown in Figure D.1. Elements that do not affect the shape of the beam delivered to the detector plane, such as the beam steering mirrors, polarizers, and power stabilizer components, are not modeled. Aperture stops, except for the spatial filter pinhole, are also not shown; diffraction effects caused by the finite apertures are not considered in the model because the apertures are assumed to be sufficiently larger than the beam so that any resulting artifacts are small. Any such artifacts that do occur are compensated by the relative aperture transmittance correction.




Figure D.1 Thin lens Gaussian optics model showing the 1550 nm optical source; mirrors, polarizers, power modulator elements, and aperture stops are not shown.

The laser beams used in the high-accuracy calibration system are approximately Gaussian beams. Gaussian beams have a Gaussian intensity profile, and propagate in a predictable manner from their waist, where the wavefront is planar. Diffraction causes the beam to spread and diverge as it propagates from the waist, with smaller beams diverging more than larger beams. In classical geometrical optics analysis, diffraction is ignored, which leads to a simple but potentially inaccurate model. A Gaussian beam ray-tracing analysis leads to a much more complex but also more accurate solution. In the following sections, both geometrical and Gaussian beam models are derived and contrasted, to show why the Gaussian model is necessary. The spatial filter pinhole selection criterion is also described.

D.1 Geometrical Optics Analysis

Using a geometrical optics analysis, the lens system shown in Figure D.1 is very simple. The distance l_0 is chosen so that a geometrically collimated beam propagates from Lens 0. Because diffraction is ignored in geometrical optics analysis, the geometrically collimated beam has a constant size and a planar wavefront at all distances beyond the lens. When the spacing between Lens 1 and Lens 2 is such that $l_1' = f_1$, and $l_2 = f_2$, the geometrically collimated beam passes through the lenses with no change, except for its size. The geometrically collimated beam then propagates from Lens 2 to the detector plane without change.

More specifically, using geometrical optics, the diameter D_0 of the collimated beam propagating from Lens 2 is related to the diameter of the collimated beam D input to Lens 1 by the following equation, where f_1 is the focal length of Lens 1, and f_2 is the focal length of Lens 2:

$$D_0 = D \frac{f_2}{f_1}. \quad (\text{D-1})$$

The resulting magnification of the collimated beam is given by $M = D_o/D = f_2/f_1$. The diameter of the beam output by the two lenses of the spatial filter can therefore be adjusted by selecting the ratio of the focal lengths of the lenses. This equation can be used to select Lens 1 and Lens 2 from the limited selection of lenses, so that D_o is close to the desired beam diameter at the detector plane. The size of the beam in the detector plane can be further adjusted by deviating from geometrical collimation, by adjusting the distance between Lens 1 and Lens 2.

Lens 1 and 2 form a compound thick-lens system. When the thin-lens approximation is used for the elements of the compound lens, the focal length of the compound lens (f_c) is given by the following equation [24, 25], where d is the distance between the lenses:

$$\frac{1}{f_c} = \frac{1}{f_1} + \frac{1}{f_2} - \frac{d}{f_1 f_2}. \quad (\text{D-2})$$

When the beam input to the compound lens is geometrically collimated, the beam output by the compound lens is also geometrically collimated when $d = f_1 + f_2$, and $f_c = \infty$. The size of the geometrically collimated beam does not change as the beam propagates away from Lens 2, so the beam diameter in the detector plane is D_o . Increasing the lens spacing d brings the focal length of the compound lens in from minus infinity, so the beam converges towards the detector plane and has a smaller diameter there. As d is increased farther, the minimum beam diameter occurs when the beam is focused on the detector plane. The beam diameter in the detector plane can therefore be controlled by adjusting the lens spacing d , at the expense of the beam's collimation. In practice, d is not increased beyond the point where the beam focuses on the detector plane. To obtain a larger beam, d is decreased from the collimated length, the focal point comes in from plus infinity, so the beam diverges towards the detector plane, and the resulting beam diameter in the detector plane is larger.

While this geometrical analysis accurately describes the mechanisms used in the optical system, diffraction causes the actual detector plane beam diameter and radius of curvature to be different from those predicted by the geometrical. Therefore a more complex model which accounts for diffraction is needed.

D.2 Gaussian Beam Ray-Tracing Analysis

Since the actual optical system uses approximately Gaussian beams, an analysis using Gaussian beams that includes diffraction affects is more accurate than the geometrical analysis. A ray-tracing analysis using the ABCD law for Gaussian beams [26] is used to model the system. The model predicts the behavior of the real system with good accuracy, because the laser beams used in the system are a good approximation of a true Gaussian beam. Measurements performed with the actual system confirm that the Gaussian beam analysis predicts the system's true performance to within a few percent.

A Gaussian beam is defined by its waist, where the beam diameter is a local minimum, and the radius of curvature of the wave front is infinite (the wave front is planar). As the beam propagates away from the waist, diffraction causes the beam to spread and diverge. The way the beam propagates from the waist is known, so only the waist location, size, and direction of propagation are needed to completely model the beam. A true Gaussian beam has a Gaussian intensity profile everywhere along its length, with radial symmetry around its propagation axis. The ray-tracing analysis follows the $1/e^2$ intensity point of the beam profile; the $1/e^2$ intensity radius (w_0) is used to define the beam's width at the waist.

Each lens in the system is analyzed separately as a region or cell, and each region is represented by its own transform matrix. Transform matrixes for the separate regions are combined by multiplying the

matrixes together. The beam travels down the z -axis, which is zero at the input to the lens in the last region. The transform matrix is accurate for $z \geq 0$ (in the region at and past the lens); the matrix for the previous region is used to predict the behavior before the lens.

To simplify the analysis, the thin-lens approximation is used. The thin-lens approximation is adequate because the distances between the principal planes of the thick lenses is small compared to the lens spacing, and the beam spreading caused by diffraction within the lenses is small compared to the beam diameter. An analysis that includes the thick lenses can be performed using dedicated optical analysis software, but details about the doublet lens's type of glasses and radiuses of curvature would be required.

Each lens in the optical system and the space around it is considered separately, as a cell. Each lens is presented with an object, which is the beam waist produced by the optical source, or the image waist produced by the previous cell. Each lens projects an image of its object waist. Positive lenses are used, so the image waist is usually a real image, projected on the side of the lens opposite the object. The ray-tracing analysis starts at the optical source's beam waist, where the beam parameter q is given by $1/q = -j(\lambda_0/\pi n w_0^2)$, where λ_0 is the vacuum wavelength of the light, and n is the index of refraction of the surrounding media. The lenses are in air, so $n = 1.0003$ is sufficient for these calculations, though a value of 1.0002 is more accurate for the actual laboratory at high altitude.

The lenses and aperture stops must have a clear aperture large enough to pass the Gaussian beam with negligible clipping, or else the Gaussian beam will be distorted. A rule of thumb is to use components with a clear aperture diameter at least 1.5 times the $1/e^2$ beam diameter at the component, so as to intercept at least 99 % of the incoming intensity [26]. In the real system, a more conservative factor is used to reduce any degradation of the approximately Gaussian beam, so aperture diameters at least three times the beam diameter are preferred. One exception is the final aperture, which is small enough to block the stray beam produced by the system's polarizer. Other exceptions are the fixed apertures on the power stabilizer components.

The lenses used in the calibration system have a clear aperture of over 20 mm, and so their diameters are 10 times that of the 2 mm diameter beam typically encountered at the lenses. However, the fixed apertures on the LPC modulator and monitor may cause problems with larger beams, since their input aperture diameter is 4 mm, and their output aperture diameter is 6 mm. The largest beam that can be passed through the optical system is limited by these apertures.

For each cell m , the beam propagates through air a distance of l_m from the object beam waist to the thin lens, which has focal length f_m . The object beam waist has a $1/e^2$ beam radius of w_{0m} . After the lens, the beam propagates through air an arbitrary distance z . Transmission matrices for the three regions in the cell, the object space, the thin lens, and the image space, are combined to form a transmission matrix for the cell. For the cell transmission matrix, $z = 0$ occurs at the input to the last region, which is the output of the thin lens in this case. The beam size and radius of curvature at an arbitrary distance past the thin lens are derived from the cell transmission matrix, using the ABCD law for Gaussian beams. It is convenient to define a characteristic length z_{0m} for each cell m , where

$$z_{0m} = \pi n w_{0m}^2 / \lambda_0. \quad (D-3)$$

The analysis gives the following general solutions for the size ($w_m(z)$) of the transmitted beam at a distance of z from lens m :

$$w_m(z)^2 = w_{0m}^2 \frac{\left(1 - \frac{z}{f_m}\right)^2 + \frac{1}{z_{0m}^2} \left[z + l_m \left(1 - \frac{z}{f_m}\right)\right]^2}{\left(1 - \frac{z}{f_m}\right) \left(1 - \frac{l_m}{f_m}\right) + \frac{1}{f_m} \left[z + l_m \left(1 - \frac{z}{f_m}\right)\right]}. \quad (D-4)$$

The radius of curvature of the wave front ($R_m(z)$) at a distance of z from lens m is given by

$$R_m(z) = \frac{\left(1 - \frac{z}{f_m}\right)^2 + \frac{1}{z_{0m}^2} \left[z + l_m \left(1 - \frac{z}{f_m}\right)\right]^2}{\frac{1}{z_{0m}^2} \left[z + l_m \left(1 - \frac{z}{f_m}\right)\right] \left(1 - \frac{l_m}{f_m}\right) - \frac{1}{f_m} \left(1 - \frac{z}{f_m}\right)}. \quad (D-5)$$

The image waist occurs a distance of l'_m from the lens, on the opposite side of the lens from the object. A negative l'_m means that a virtual image waist occurs on the same side of the lens as the object waist. The location of the image waist is determined by solving for where the radius of curvature is infinite ($1/R_m(z) = 0$). The resulting image distance is given by

$$l'_m = f_m \frac{l_m^2 - l_m f_m + z_{0m}^2}{l_m^2 - 2l_m f_m + f_m^2 + z_{0m}^2}. \quad (D-6)$$

Equation (D-6) can be expressed in the following form, which is similar to the classic *lens equation* [25], derived using geometrical optics analysis. This equation may therefore be considered a *modified lens equation for Gaussian beams* [27]:

$$\frac{1}{f_m} = \frac{1}{l'_m} + \frac{1}{l_m + \left(\frac{z_{0m}^2}{l_m - f_m}\right)}. \quad (D-7)$$

The sign of the second term differs from that in the referenced lens equation because of differences in the sign conventions used. Here, the distance from the thin lens to the object (l_m) is considered positive, while in the reference the distance is considered negative; the difference results in the sign change in the referenced equation. Also, since the thin lens is in air, $f'_m = f_m$.

Equation (D-7) shows that when $|l_m - f_m| \gg z_{0m}^2$, the system conforms to the geometrical optics analysis, and the Gaussian beam waist is imaged like any other object. But in the region around $l_m = f_m$, the geometrical optics solution has appreciable error. All the lenses in the calibration system are used with $l_m \approx f_m$, so the Gaussian beam optics solution is needed to obtain accurate results. It should be noted that when the above equations are numerically evaluated, the semi quadratic form shown in eq (D-6) is preferable to the form shown in eq (D-7) because the latter has a division by zero when $l_m = f_m$, which can cause problems in the numerical computations and can cause excessive roundoff error when l_m is close to f_m , and $1/(l_m - f_m)$ is very large.

Several interesting results are obtained from the Gaussian beam analysis, perhaps the most significant of which is how a Gaussian beam is collimated. A geometrically collimated beam is usually thought of as having an infinite radius of curvature, or planar wave front and finite extent at all distances. The Gaussian beam waist is the only place in a Gaussian beam which meets this criterion if the requirement of finite extent is neglected. So in a sense, all Gaussian beams are collimated, at least at their waist. Around the waist, the beam is considered collimated within the Rayleigh range of the waist [28], which is given by $z_R = z_{0m}$. Since all Gaussian beams are collimated near their waist, of what distance should the lenses of the spatial filter be spaced to collimate the transmitted Gaussian beam?

The transmitted Gaussian beam is considered optimally collimated when it has the lowest divergence in the far field, which occurs when $l_m = f_m$. When $l_m = f_m$, the object beam waist is placed at the primary focal point of lens m , where the point source is placed in geometrical collimation. With such an optimally collimated Gaussian beam, the image waist is projected to the lens's secondary focal point, and not to infinity as predicted by the geometrical analysis. This condition results in the largest image beam waist, which diverges the least since large Gaussian beam waists diverge more slowly than smaller waists, so in the far field the smallest diameter beam and smallest divergence angle is produced when the waist is largest. The largest beam waist also produces the longest Rayleigh range. However, in the real system, since the detector plane is not very far from the lenses, a smaller diameter detector plane beam can usually be produced when Lens 2 is in a position slightly different from the optimal collimation location, similar to the way the focused beam is smaller than the collimated beam in the geometrical optics analysis. The size of the optimally collimated beam waist image is given by $w_{cm} = w_m(f_m)$, where $l_m = f_m$, or

$$w_{cm} = w_{0m} \frac{f_m}{z_{0m}} = \frac{\lambda_0 f_m}{\pi w_{0m}}. \quad (D-8)$$

Another interesting result of the Gaussian beam analysis is how under certain circumstances (when $|l_m - f_m| \gg z_{0m}^2$), the Gaussian beam waist is imaged in accordance with geometric optics. In this case, the lenses in the spatial filter can be considered as a compound lens with a variable focal length, as previously described, that images the laser source's waist. Then the image and object distances are related by the focal length of the compound lens, in accordance with the classic lens equation, and the image distance can be increased by increasing the object distance. This result is significant in that it shows how in the calibration system, the final image beam waist can be projected farther from Lens 2 by moving the laser source farther from Lens 1. Therefore by moving the laser source farther away, it is possible to project the image waist all the way to the detector plane and beyond. If the image waist is projected to the detector plane, the detector will be illuminated with a true planar wave front. This effect is demonstrated later in this section; the analysis uses the Gaussian beam result, because using the geometrical optics analysis on the compound lens is greatly complicated by the shift in the location of the principle planes of the compound lens when the focal length is adjusted (the compound lens cannot be considered "thin" when the laser and detector plane are at reasonable distances). However, understanding the relationship to geometric optics does provide some insight into the seemingly bizarre behavior of the Gaussian beam.

Another interesting result of this analysis is that a Gaussian beam waist image can never be projected to infinity. Doing so would require an infinite object waist distance, in which case diffraction would also result in an infinitely large beam. Since an infinite object distance cannot be achieved, this condition cannot be produced in the laboratory.

The high-accuracy calibration system's optical system is shown with Gaussian beams in Figure D.1. In the figure, the rays shown follow the $1/e^2$ intensity points of the Gaussian beam. The figure shows all the

lenses in the system, the image and object distances (l_m and l'_m) for each cell m , and the distances d_1 (the distance between the laser's output and Lens 1) and d_2 (the distance from Lens 2 to the detector plane). The distance d_1 is different for each laser, but d_2 is considered to be constant even though it changes slightly when Lens 2's position is adjusted, because d_2 is much greater than the amount of adjustment. The mirrors, polarizers, and other elements of the optical system are not shown in the figure, because they do not significantly affect the Gaussian beam.

The 1550 nm optical fiber coupled laser is depicted in Figure D.1; this laser has a single-mode, polarization-maintaining optical fiber pigtail that is permanently attached to the laser's gain medium. The single-mode optical fiber core produces a very nearly Gaussian profile at the end of the optical fiber where the core is cleaved [29]. The mode field diameter of the single-mode optical fiber core is used as the Gaussian beam waist diameter [8]. A collimating lens (Lens 0 in Figure D.1) is used to generate an optimally collimated, free-space beam from the optical fiber. Optimal collimation is achieved by positioning Lens 0 so that the end of the optical fiber's core is at the lens's primary focal point; then the image waist is formed at the secondary focal point.

The other lasers used in the system have a free-space output with a single transverse mode, TEM_{00} , which has a flux density that is ideally Gaussian over the beam's cross section [30]. With these lasers, the beam waist generated inside the laser's resonant cavity is analogous to the waist in the optical fiber core, and the laser's collimating optic is analogous to Lens 0. The fiber-coupled laser is used in this analysis because its beam waist diameter and collimating lens focal length are known, where they are unknown for the free-space lasers. The free-space lasers are analyzed with the same technique, by estimating the values for cell 0. For these lasers, the output beam diameter specified by the manufacturer is assumed to be the $1/e^2$ intensity diameter of the waist image produced by the laser's collimating optic. And the distance l'_0 is considered negligible compared to d_1 so $l_1 \approx d_1$, and the waist is assumed to occur at the laser's output. The analysis begins in cell 1, using one-half of the manufacturer's specified beam diameter as w_{01} . The beam diameter at Lens 1 (D in Figure D.1) is measured, and the values of l_1 and/or w_{01} are adjusted so that the theoretical value of D is consistent with the measured value of D . The discrepancy between the theoretical and measured values of D is due in part to imperfections in the approximately Gaussian beam, as well as to the assumptions made about w_{00} and l'_0 , so the adjustment is usually necessary even with optical fiber-coupled sources.

The parameters of interest for each cell in Figure D.1 are determined as follows. In cell 0, the initial Gaussian beam waist (w_{00}) is produced at the end of the optical fiber, and has width determined from the optical fiber's mode-field diameter, which for the 1550 nm laser is 9.5 μm . The equivalent Gaussian beam waist radius is thus 4.75 μm , (one-half of the $1/e^2$ diameter of the mode field). Lens 0 is an anti-reflection coated glass aspherical lens with focal length of 6.2 mm, so $f_0 = 6.2$ mm. The collimator is assumed to optimally collimate the beam, so $l_0 = l'_0 = f_0$. The image beam waist's size is $w_{01} = w_{c0}$, from eq (D-8) with $m = 0$, and w_{01} is the input to cell 1. The diameter at the input of Lens 1 (D) can be calculated using the matrix for cell 0 with a distance of $z = d_1$, so $D = 2w_0(d_1)$, which is equivalent to the matrix for cell 1 with $z = 0$ or $D = 2w_1(0)$. In cell 1, the object waist distance is given by: $l_1 = d_1 - l'_0$, and the image distance (l'_1) is given by eq (D-6) with $m = 1$. The size of the image waist projected by Lens 1 is given by $w_{02} = w_1(l'_1) = w_{c1}$. The spatial filter pinhole is placed around w_{02} , where it primarily limits the field of view of the system; ideally, the Gaussian beam is not affected by the pinhole, as previously described. The size of the spatial filter pinhole does not affect the spatial filter's numerical aperture or f-number, which is a function of the lens diameter and focal length.

Lens 2 is initially placed so that the transmitted beam is optimally collimated, by placing the lens one focal length away from the spatial filter pinhole, or minimizing the spot size in the far field. Precise positioning of Lens 2 is not necessary at this point in the alignment, so approximate optimal collimation is

sufficient. When optimally collimated, the object and image waist distances are equal to the focal length of the lens, so $l_2 = l'_2 = f_2$. In the detector plane, the beam size (w_{dp}) and radius of curvature (R_{dp}) are given by eq (D-10) and eq (D-11) with $m = 2$ and $z = d_2$. However, the system is seldom used with the output beam optimally collimated; in practice the detector plane beam size is adjusted by changing the distance between Lens 2 and w_{02} , which is l_2 .

The result of adjusting l_2 on the beam size and radius of curvature in the detector plane (w_{dp} and R_{dp}), can be determined by substituting, in eq (D-4) and eq (D-5), l_2 with $f_2 + \delta$, and z with d_2 . Then, δ is the amount Lens 2 is moved from the optimally collimated position. The result of the substitution on w_{dp} and R_{dp} is shown in Figures D.2 and D.3, using real values for the parameters: $\lambda_0 = 1550$ nm, $n = 1.00027$, $w_{00} = 4.75$ μ m, $f_0 = l_0 = 6.2$ m, $d_1 = 1.04$ m, $f_1 = f_2 = 0.1$ m, $l_2 = f_2 + \delta$, and $d_2 = 1.28$ m. While these curves can change significantly with different sources or lenses, all of the lasers currently used in the system have a similar relationship between δ and w_{dp} , which is convenient for producing the 2 mm diameter beam in the detector plane ($w_{dp} = 1$ mm). To date, a detector plane beam size of only a 2 mm has been used with the system.

Figure D.2 shows that a nearly linear relationship between δ and w_{dp} results. However, Figure D.3 shows that if the detector plane beam size is reduced to its minimum, a shorter radius of curvature results. How the projected waist image distance (l'_2) varies with δ is shown in Figure D.4, using eq (D-6) with $m = 2$ and $l_2 = f_2 + \delta$. The largest waist is projected to $l'_2 = f_2$ when $\delta = 0$ (the beam is optimally collimated). The smallest w_{dp} is obtained with a positive δ , which projects a smaller beam waist closer to the detector plane. Increasing δ farther makes the image waist even smaller, but moves it back towards the lens, resulting in an increased detector-plane beam diameter, but with a shorter radius of curvature. So to increase the detector-plane beam size, a negative δ is used, which causes a smaller waist to be projected farther from the detector plane, and the smaller waist diverges more, resulting in a larger beam in the detector plane. A sufficiently large negative δ causes the waist image to be projected to the left of Lens 2, resulting in a virtual waist image. However with a negative δ , the resulting detector plane beam has a

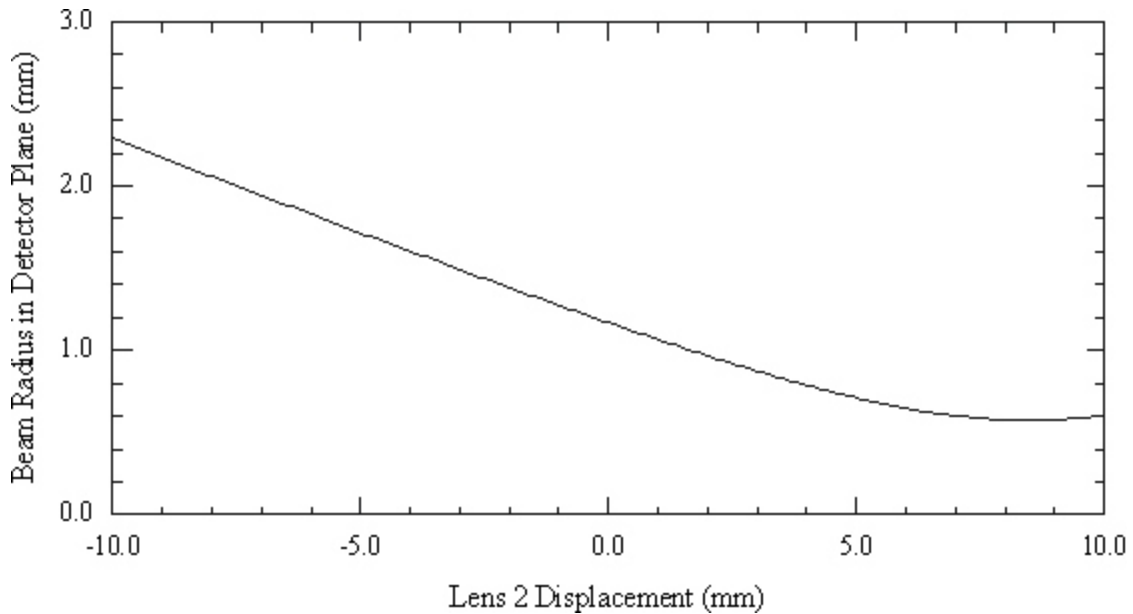


Figure D.2 The size (radius) of the Gaussian beam in the detector plane versus the displacement of Lens 2 from the optimally collimated location.




Figure D.3 The radius of curvature of the wave front in the detector plane versus the displacement of Lens 2 from the optimally collimated location.

longer radius of curvature than it would with a large positive δ , so using a negative δ is the preferred method for expanding the beam.

Using an increasingly negative δ will further expand the size of the detector plane beam, but at the expense of a shorter radius of curvature. Additionally, the physical layout of the lens and pinhole limit how negative δ can be. So if a larger beam is required, it is better to adjust other parameters in the system (such as the lens focal lengths or d_1), and keep the system near the point of optimal collimation by limiting the range of δ that is used. The fixed apertures on the LPC monitor element will ultimately limit the maximum beam size that can be supported without excessive diffraction artifacts. Similarly, if a detector plane beam size that is smaller than that achieved with a positive δ is required, other parameters must be adjusted to reduce the minimum beam size.

The longest radius of curvature shown in Figure D.3 occurs when δ is adjusted so that the image waist is projected the longest distance, as shown in Figure D.4 for the same δ . The radius of curvature in this configuration is long, because the beam waist is projected close to the detector plane. If the waist could be projected a little farther so that it occurs at the detector plane, the radius of curvature would approach infinity. Recall the result which showed that in certain circumstances, the spatial filter lens system can be considered a compound lens that produces a geometric image of the laser's waist, so by increasing the object waist distance, the image waist distance should also increase. This result is confirmed by the Gaussian beam analysis; increasing the distance between the laser's output and Lens 1 (d_1 in Figure D.1), does indeed increase the distance to the final image waist (l_2), for certain lens displacements δ . For the 1550 nm laser source, increasing d_1 by 1 m results in the detector plane radius of curvature shown in Figure D.5. In this configuration, the waist can actually be projected behind the detector plane, so it crosses the plane twice, with $\delta \approx 0.5$ mm, and $\delta \approx 8$ mm.




Figure D.4 The distance from Lens 2 to the projected waist image (l_2') versus the displacement of Lens 2 from the optimally collimated location.

The absolute value of the radius of curvature is shown in Figure D.5, so that a logarithmic axis could be used; the radius of curvature is actually negative when δ is between 0.5 and 8 mm, because the waist occurs behind the detector plane, and the rays are therefore converging as they approach the detector plane, so the radius of curvature is negative. The detector-plane variation in beam size has a form very




Figure D.5 The radius of curvature of the wave front in the detector plane versus the displacement in Lens 2 from the optimally collimated location, when the 1550 nm laser is moved 1 m farther away from Lens 1.

similar to that shown in Figure D.2, except that the beam radius varies from 0.35 to 3.5 mm. As δ is increased from zero, the first time the waist crosses the detector plane it has a diameter of approximately 3 mm. As δ is increased farther, the second time the waist crosses the detector plane its diameter is approximately 0.8 mm. A similar situation occurs when the laser is placed in its original location and a lens of 50 mm focal length is used as Lens 1, in place of the lens with 100 mm focal length.

This analysis shows that a wide range of detector plane beam diameters and radii of curvatures can be obtained by adjusting the location of Lens 2. However, if the range of adjustment is not sufficient, other parameters such as the lenses' focal lengths and distance to the laser can be adjusted. However, the result of changing the lens focal length and laser distance is complex, so an ad hoc, trial-and-error approach is used to determine which configuration adequately provides the desired beam characteristics. Dedicated mathematics software is used to evaluate the Gaussian beam equations that predict the behavior of the system, so that several different configurations can be tested virtually before the actual system is aligned. After the desired configuration is chosen, the system is aligned and δ adjusted to achieve the desired beam diameter in the detector plane. If the beam waist is being projected to the detector plane, δ is adjusted to position the location of the beam waist, instead of the beam size.

The actual lens displacement δ that is used in the system is determined experimentally by measuring the detector plane beam diameter while δ is adjusted. The theoretical value for δ that is used in the final Gaussian beam analysis is determined using the measured detector-plane beam diameter, so the theoretical δ is consistent with the average measured detector plane beam diameter. Imperfections in the approximately Gaussian beam and lenses cause the actual system to behave slightly differently from the theoretical system, but measurements performed to date show that the actual system usually performs to within a few percent of the theoretical system.

D.3 Gaussian Beam Parameter Specification

Several parameters describing the Gaussian beam are provided in the calibration report. The incidence angle of the beam, relative to the detector's faceplate or entrance aperture is measured and reported. The laser beam's incidence angle is described as the difference from normal incidence. The incidence angle is measured by holding a glass flat to the detector's entrance aperture or faceplate, and locating the beam reflected from the glass flat. The difference between the incidence angle and normal incidence (θ_I) is calculated by measuring the distance between the incident and reflected beam (d_I), at a distance (D_I) in front of the detector, using the equation:

$$\theta_I = \tan^{-1}\left(\frac{d_I}{2D_I}\right) \approx \frac{d_I}{2D_I}. \quad (\text{D-9})$$

The approximation is accurate when θ_I is small. The incidence angle is minimized during detector alignment, so is usually very small (the incidence angle is usually a fraction of a milliradian).

The full-width beam divergence angle at the detector's entrance aperture is specified in the calibration report to allow for simple geometrical modeling. To support a better Gaussian optics model, the distance from the detector plane to the final beam waist (d_{dp} in Figure D.1), and the final beam waist radius (w_{03} in Figure D.1) are also specified in the report. These parameters are determined by the Gaussian beam model using the theoretical lens displacement δ , which is consistent with the average measured detector plane beam diameter. Again, measurements have shown that the actual parameters usually agree within a few percent with the theoretical parameters.

The full-width beam divergence angle in the detector plane (θ_{dp}) can be derived from the radius of curvature in the detector plane (R_{dp}), given by eq (D-5) with $m = 2$, $l_2 = f_2 + \delta$, and $z = d_2$. The equation relating θ_{dp} to R_{dp} is given by the following; the approximation is accurate when θ_{dp} is small:

$$\theta_{dp} = 2 \tan^{-1} \left(\frac{w_{dp}}{R_{dp}} \right) \approx 2 \frac{w_{dp}}{R_{dp}}. \quad (D-10)$$

When the beam waist occurs sufficiently far from the detector plane ($d_{dp} \gg z_{03}$), the detector plane beam divergence can be measured directly, by measuring the beam size in the detector plane (w_{dp}), and at a location a short distance d beyond the detector plane (w_{dp+d}). The detector plane beam divergence is then given by

$$\theta_{dp} = 2 \tan^{-1} \left(\frac{w_{dp+d} - w_{dp}}{d} \right) \approx 2 \frac{w_{dp+d} - w_{dp}}{d}. \quad (D-11)$$

The approximation is accurate for small θ_{dp} . Similarly, the radius of curvature can be determined directly from the measured diameters:

$$R_{dp} = d \frac{w_{dp}}{w_{dp+d} - w_{dp}}. \quad (D-12)$$

However, the analysis used to derive eq (D-11) and eq (D-12) is based on the assumption that the divergence angle is constant for $d_2 < z < d_2 + d$. The equations are therefore accurate in the far-field of the beam waist, where the divergence angle is approximately constant. Nearer to the beam waist the divergence is changing rapidly. So unless d is small compared to the Rayleigh range of the waist (z_{03}), the result can be erroneous. Therefore when the waist occurs near the detector plane, the beam divergence from the Gaussian analysis is usually more accurate.

When the detector plane is sufficiently into the far-field of the beam waist, the result for divergence angle from either eq (D-10) or eq (D-11) is reported; they typically differ by at most only a few percent. But the waist size (w_{03}) and distance (d_{dp}) that are reported are the result of using the theoretical δ in eq (D-4) and eq (D-5). This is the mode in which the system has been used to date. However in the future the beam waist may be projected closer to or behind the detector plane, as previously described. In these cases, instead of adjusting the position of Lens 2 to obtain the desired detector plane beam size, the lens position is adjusted to obtain the desired waist location. The waist can be located by moving the beam size measurement device along the z-axis until the local minimum is located. When the waist location coincides with the location of the detector plane, the measured detector plane beam size is also the waist size ($w_{03} = w_{dp}$), the divergence angle is zero ($\theta_{dp} = 0$), and the distance from the beam waist to the detector plane is zero ($d_{dp} = 0$).

D.4 Spatial Filter Pinhole Selection

Lens 1 and 2 form a classic spatial filter [7], where the spatial characteristics of the incident beam are broken up into their Fourier components in the transform plane. A pinhole aperture is placed in the center of the transform plane, where it ideally passes only the spatial mode of lowest order. However, when used with coherent light sources, diffraction around the edge of the pinhole causes undesired artifacts,

which complicates the selection of the pinhole size. In the high-accuracy calibration system, a relatively large pinhole is used to minimize the degradation of the laser beam, already only approximately Gaussian.

The calibration system's spatial filter consists of Lens 1, Lens 2, and the pinhole mount with precision pinhole installed. Both lenses are positive, so using geometrical optics analysis, the pinhole is placed at the secondary focal point of Lens 1, where a plane wave incident on the lens is focused. The pinhole itself is a circular aperture. Lens 2 is placed so that the focused spot and pinhole are at its primary focal point, so that the image of the focused spot is projected to infinity. This arrangement can be used with lenses of dissimilar focal lengths to expand or contract the beam; in common beam-expander devices for example, a microscope objective is used as Lens 1, with a much longer focal length Lens 2. But in the high-accuracy calibration system, the focal lengths of both lenses is the same or similar, because no magnification is desired. Lenses of relatively long focal length are used to provide sufficient working distances.

The image formed in the focal plane of Lens 1 is the two-dimensional Fourier transform of the spatial frequencies present in the beam that is incident on the lens. The secondary focal plane of Lens 1 is therefore called the transform plane. The pinhole is placed in the transform plane, centered around the secondary focal point of the lens. Only the lowest spatial frequencies pass through the pinhole; higher frequencies are blocked. Lens 2 provides the inverse transform, which reconstructs only the lowest frequencies of the incident beam. The arrangement therefore forms a spatial low-pass filter. While the concept is simple, selecting the best diameter of pinhole to use is more complex.

Diffraction of the coherent light at the edge of the pinhole distorts the reconstructed beam. Truncation of the Gaussian beam by finite apertures, such as the spatial filter pinhole, results in diffraction artifacts that resemble Fresnel zone plates, being superimposed on top of the ideal Gaussian profile [28, 31]. Therefore, selection of the spatial filter pinhole size is a trade-off between effective spatial filtering, and minimizing the diffraction artifacts produced by the pinhole. As a compromise, a pinhole is used with diameter that is at least 2.3 times the $1/e^2$ intensity diameter of the Gaussian beam waist that occurs in the spatial filter. Then, approximately 100 % of the incident power is transmitted, and the ripples caused by diffraction have a maximum amplitude of approximately 1 % [28]. Therefore, the minimum pinhole diameter D_p is given by the following, where w_{02} is the radius of the Gaussian beam waist that is formed in the spatial filter:

$$D_p = 4.6w_{02}. \quad (D-13)$$

The Gaussian beam waist radius w_{02} can be determined using several methods. The least uncertain method is direct measurement of the beam waist diameter using the scanning-slit beam profiler; then imperfections in the Gaussian beam are included in the width measurement. But if the beam waist is too small or the power density too high for the beam profiler, the theoretical waist size from the Gaussian beam analysis is used. The following equation is commonly used to estimate the focused Gaussian beam waist radius [26, 32, 33], where D is the measured $1/e^2$ intensity diameter at Lens 1. In this equation, the incident beam waist is assumed to occur at Lens 1, where D is measured; since in the real system the waist does not occur at the lens, the relationship is approximate:

$$w_{02} \approx 2 \frac{f_1 \lambda}{\pi D} \approx 0.637 \frac{f_1 \lambda}{D}. \quad (D-14)$$

In practice, the spatial filter pinhole is selected from a small number of available pinhole sizes, with diameters typically spaced 10 to 100 μm apart. Selecting too large a pinhole is better than selecting too small a pinhole, since we are more concerned with minimizing diffraction artifacts than with filtering the already approximately Gaussian beam. The calculated diameter D_p is used as a minimum, so the smallest pinhole that is at least as large as D_p is used, unless D_p is only slightly larger than a smaller pinhole, in which case the smaller pinhole is used. Pinholes from various manufacturers are used, but all have a precision pinhole cut in a thin metal substrate. The pinholes are visually inspected for damage before use, but the power levels usually used in the calibration system will not damage the pinholes, even when the beam is focused.

The pinhole diameter D_p from eq (D-13) is somewhat larger than that recommended by most spatial filter manufacturers, but selection of a pinhole potentially larger than necessary has other benefits. Selecting a larger pinhole provides more room for error in the estimate of w_{02} . Also, pointing instability in the laser can make the location of the focused spot change slightly with time; so using a larger pinhole helps prevent the drift from causing a severe error in pinhole alignment. The larger pinhole is also easier to align, with more tolerance for alignment error. To minimize the alignment error, a precision pinhole mount with differential micrometers is used.

When precision beam profile measurements with sufficient dynamic range are available, the actual beam profile in the detector plane can be measured, and the effect of different spatial filter pinhole diameters better assessed. Other filtering techniques may also be explored. A scanning pinhole will likely be used for the precision measurements of profile, but a good camera system may also be usable over a limited range of wavelengths.

APPENDIX E. Relative Aperture Transmittance: Estimation and Uncertainty

A Gaussian laser beam can result in power meters with different sized or shaped limiting apertures to absorb different amounts of power, because of the difference in their aperture's transmittance. To minimize the discrepancy, detectors with limiting apertures that are large compared to the beam size are usually used. For example, a circular limiting aperture with a diameter that is twice the beam's $1/e^2$ intensity diameter will transmit 99.97 % of the total power in a true Gaussian beam. Since a loss of 0.03 % of the incident power is significant in high-accuracy calibrations, a larger relative aperture size is usually used; a circular aperture three times the size of the beam diameter transmits all of the incident power in an ideal Gaussian beam, to within a fraction of a part per million. Unfortunately, the real beam is only approximately Gaussian, and much more power is blocked by the aperture than predicted by the ideal Gaussian model. With the actual laser beams used in the high-accuracy calibration system, a measurable power loss typically occurs even with circular apertures that have a diameter that is four times the $1/e^2$ beam diameter. To compensate for any such discrepancy in the power absorbed by the standard and test detectors, the relative aperture transmittance correction is applied to the test detector calibration result.

Determining the relative aperture transmittance for a test detector requires knowledge of the test detector's apparent aperture. When the shape of the detector's limiting aperture is in question, a spatial uniformity scan of the test detector is necessary to determine the detector's apparent aperture. The aperture is in question when the detector does not have an obvious physical limiting aperture. If a spatial uniformity scan is required but the customer has not ordered a formal scan (which is a separate special test), a quick scan with sufficient resolution for determination of the detector's apparent aperture is performed as part of the high-accuracy calibration. When such a scan is needed, a figure showing the detector's apparent aperture is included in the report, but the actual uniformity scan data is not provided. The apparent aperture of a detector with no limiting aperture is shown in Figure 6 (of the main document).

When NIST is able to perform sufficiently precise measurements of beam profile, the beam profile itself can be supplied to the end user. However, the beam profile cannot currently be measured with sufficient precision to accurately estimate the small correction for relative aperture transmittance. Currently a measurement technique that uses fixed, circular apertures is used to estimate the relative transmittance. In the future, a measurement technique using a scanning detector with a pinhole aperture could be used to measure the beam profile with sufficient precision for the measured profile and the detector's spatial uniformity to be combined to calculate the relative transmittance without requiring the measurements of circular aperture transmittance.

The additional power present in the real beam results from imperfections in the optical system that include diffraction artifacts produced by the finite apertures in the optical system and laser source, light scattered by the air and the optical elements, aberration caused by imperfect lenses, and multiple beams caused by multiple reflections in optical elements such as the polarizers. And in addition to the diffraction caused by its finite apertures, imperfections in the laser's internal alignment can distort the beam at the source. Even with optical fiber-coupled lasers, the beam produced at the end of the fiber core is not purely Gaussian, but can contain cladding modes and scattered power [8]. Because of these shortcomings, there is enough power away from the center of the beam to cause problems at the uncertainty levels of interest in the high-accuracy calibration system.

The aperture referred to in this discussion is the detector's apparent aperture. For detectors that have true limiting apertures, the apparent aperture is the same as the limiting aperture. However, not all detectors have physical limiting apertures. For detectors that have no physical limiting aperture, the apparent

aperture estimates an equivalent limiting aperture. With these detectors, the apparent aperture may not be exactly equivalent to a limiting aperture, because with a true limiting aperture, the detector has no response to light outside the boundary defined by the aperture. But when there is no true limiting aperture the detector may have some response to light outside its apparent aperture. Also, even if a detector does have a true limiting aperture, its responsivity may not be uniform within the boundaries of its apparent aperture. The relative aperture transmittance for these detectors is measured in such a way that helps accounts for these effects, as described later.

E.1 Measured and Theoretical Transmittance for Circular Apertures

The calibration system's beam has significantly more power away from its center than a true Gaussian beam would, so a correction based on calculations using the theoretical Gaussian profile is not accurate; therefore measured results are needed. When the calculated relative aperture transmittance for the theoretical Gaussian beam is compared to the measured results for the real beam, the difference is clearly evident. The relative aperture transmittance for a true Gaussian is derived first, for comparison with the measured results.

A Gaussian profile beam has an intensity profile $I_G(r)$ given by the following equation [32], where w is the $1/e^2$ beam radius, r is the distance from the center of the beam, and A is an arbitrary scale factor:

$$I_G(r) = A \exp(-2r^2 / w^2). \quad (E-1)$$

The amount of power P that falls within a circular aperture with radius R , centered around the beam, is given by

$$P(R) = \int_0^R I_G(r) r dr. \quad (E-2)$$

When eq (E-2) is solved analytically, the following is obtained [28]:

$$P(R) = A \frac{w^2}{4} \left(1 - \exp\left(-2 \frac{R^2}{w^2}\right) \right). \quad (E-3)$$

The relative aperture transmittance is expressed as $T_{RA}(R, C)$, which means the transmittance of aperture R relative to the transmittance of aperture C . For centered, circular apertures, $T_{RA}(R, C)$ is the ratio of the power transmitted by a circular aperture with radius R to the amount of power transmitted by the circular reference aperture with radius C :

$$T_{RA}(R, C) = \frac{P(R)}{P(C)}. \quad (E-4)$$

When eq (E-3) is combined with eq (E-4) and simplified, the following relationship for the relative aperture transmittance of the theoretical Gaussian beam, in centered circular apertures, is obtained:

$$T_{RA}(R, C) = \left(\exp\left(-2 \frac{R^2}{w^2}\right) - 1 \right) / \left(\exp\left(-2 \frac{C^2}{w^2}\right) - 1 \right). \quad (E-5)$$

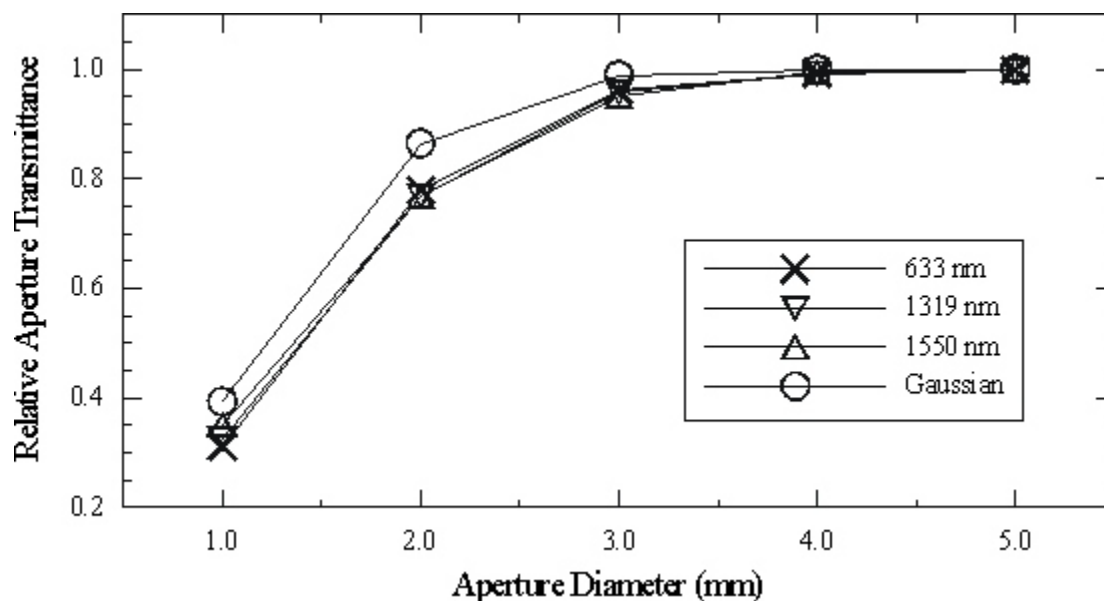


Figure E.1 Measured and theoretical relative aperture transmittance for small circular apertures, relative to a circular reference aperture of 10 mm diameter, using a beam of 2 mm diameter.

In the high-accuracy calibration system, the relative aperture transmittance is measured. Figures E.1 and E.2 show measured results along with the calculated result from eq (E-5) for the true Gaussian beam. Figure E.1 shows the results for apertures 1 to 5 mm in diameter, Figure E.2 shows the results for apertures 5 to 10 mm in diameter. The results are also listed in Table E.1, where the standard uncertainty of the measurements is also given. The transmittance shown is for a beam of nominally 2 mm diameter in a centered, circular aperture, and is relative to a circular reference aperture of 10 mm diameter. Figure E.1 shows the relative aperture transmittance for apertures with diameters from 1 to 5 mm. In




Figure E.2 Measured and theoretical relative aperture transmittance for large circular apertures, relative to a circular reference aperture of 10 mm diameter, using a beam of 2 mm diameter.

practice a smaller beam would probably be used with such small detectors, but the figure clearly shows the significant difference between the measured and true Gaussian beam transmittance. Figure E.2 shows the relative aperture transmittance for more practical aperture sizes, where essentially all of a true Gaussian beam would be transmitted; again the difference between the measured and Gaussian beam results is clearly evident. The difference between the measured transmittance for the different wavelengths is usually small, but the transmittance of the 1319 nm wavelength is exceptionally low, possibly due to increased scattering and distortion caused by the two narrow-band interference filters that are used with the laser. The two 633 nm measurements show that the relative aperture transmittance can change slightly when the system is realigned. When a sufficient history of measurements is available, it may be possible to use the average of previous transmittance measurements instead of measuring the transmittance each time the system is used. However for quality control purposes, even when a transmittance history is available, it is likely that some measurements will be performed each time the optics are re-aligned, to ensure that the current alignment is comparable to the historic alignments.

Not shown in the measured data is the relative aperture transmittance for apertures that are larger than the reference aperture. In this case, the relative aperture transmittance can be >1, because the larger aperture can transmit more power than can the reference.

Table E.1 Theoretical and measured relative aperture transmittance for circular apertures and a 2 mm diameter beam. The standard uncertainty of the measured results is shown.

Aperture Diameter (mm)	Laser Beam (2 mm 1/e ² intensity diameter)						
	Gaussian	1550 nm*	1319 nm*	633 nm*	633 nm #2	515 nm	488 nm
1	0.393469	0.3519 ±0.0013	0.3247 ±0.0011	0.31272 ±0.00021	N/A	N/A	N/A
2	0.864665	0.7693 ±0.0014	0.7678 ±0.0014	0.77806 ±0.00044	N/A	N/A	N/A
3	0.988891	0.95012 ±0.00089	0.95670 ±0.00087	0.96020 ±0.00015	N/A	N/A	N/A
4	0.999665	0.99424 ±0.00046	0.99206 ±0.00046	0.992424 ±0.000095	N/A	0.9984 ±0.0020	N/A
5	0.999996	0.99880 ±0.00025	0.99755 ±0.00024	0.998884 ±0.000053	0.99941 ±0.00033	0.99950 ±0.00063	0.99968 ±0.00030
6	1.000000	0.99952 ±0.00016	0.99921 ±0.00016	0.999809 ±0.000035	0.999852 ±0.000049	0.999892 ±0.000084	0.999897 ±0.000040
7	1.000000	0.99972 ±0.00013	0.99966 ±0.00013	0.99977 ±0.00011	0.999922 ±0.000046	0.999932 ±0.000043	0.999925 ±0.000022
8	1.000000	0.99988 ±0.00011	0.999785 ±0.000083	0.999951 ±0.000026	0.999963 ±0.000036	0.999965 ±0.000039	0.999955 ±0.000012
9	1.000000	1.000000 ±0.000062	0.999840 ±0.000070	0.999974 ±0.000023	0.999982 ±0.000050	0.999974 ±0.000032	0.999982 ±0.000011
10	1.000000	1.000000 ±0.000062	1.000000 ±0.000070	1.000000 ±0.000025	1.000000 ±0.000035	1.000000 ±0.000030	1.000000 ±0.000009

N/A - Not Available.

* - Measured using a slightly different process and uncertainty analysis.

The measured relative aperture transmittance for three additional wavelengths (488 nm, 515 nm, and 633 nm #2) are shown in Figure E.2; they were acquired more recently than the others, and better reflect the current uncertainty in the measurement. The other measurements were acquired with an older measurement process, and have a less conservative uncertainty analysis. Also, the lower uncertainty in the older measurements is also due in part to the more stable laboratory environment that was present during the earlier measurements. The transmittance of the smaller apertures was not measured in the more recent acquisitions, because of environmental instability, which necessitated additional measurements of the larger aperture's transmittance to obtain a sufficiently small uncertainty in the average. Also, the transmittance of the small apertures was not needed for the calibrations performed.

The measured data show that the magnitude of the relative aperture transmittance correction can be very significant at the uncertainty levels of interest in the high-accuracy calibration system, where the goal is a standard uncertainty of 0.01 %. For example the magnitude of the correction, $100 \% \cdot |1 - T_{RA}|$, can be calculated for the transmittance results. The average magnitude of the correction for an 8 mm test detector aperture relative to a standard with an aperture of 10 mm diameter, when using a beam of 2 mm diameter, is 0.008 %, which is almost as large as the entire desired error budget. And the average measured correction for a test detector aperture of 5 mm diameter increases to 0.1 %, about 10 times the desired error budget! Therefore, ignoring the relative transmittance can cause a significant and unaccounted error in high-accuracy calibrations, whenever the standard and test detector have dissimilar apertures.

Even when the theoretical Gaussian beam is considered, the error that results from ignoring the relative aperture transmittance can be prohibitive. As shown by the data for the theoretical Gaussian in Table 1, using a circular aperture that is 1.5 times the Gaussian beam diameter (a 3 mm diameter aperture with a 2 mm diameter beam) would result in a loss of 1.1 % of the power that is delivered through a 10 mm diameter aperture. An aperture with a diameter that is two times the beam diameter will miss 0.035 % of the power, which is again greater than the desired standard uncertainty. And the situation with the real beam is worse; using the measured results, an aperture three times the beam diameter will typically miss 0.03 % of the power that is transmitted through a 10 mm diameter aperture.

The relative aperture transmittance for the theoretical Gaussian beam becomes insignificant (at the uncertainty levels of interest) for apertures with a diameter two and a half times the beam diameter or more. But the measured transmittance for the actual laboratory beam can still be significant even when the aperture is more than four times the beam size. The measured data in Table 1 show that the average magnitude of the correction even for a relatively large aperture of 9 mm diameter (which is four and a half times the 2 mm beam size) is 0.004 %, approximately half of the desired error budget. Therefore, the correction for relative aperture transmittance is applied to all detectors, even those with relatively large apertures.

Since the relative aperture transmittance is always applied, the magnitude of the uncertainty in the correction can limit the accuracy of the calibration system. For the more recent measurements shown in Table 1, the standard uncertainty of the correction typically exceeds 0.01 % for apertures smaller than three times the beam diameter. Therefore, test detectors with an apparent aperture that has a diameter greater than three times the beam's diameter are recommended. Even if both the standard and test detector have identical but small apertures, the uncertainty may still be prohibitive, primarily because of the uncertainty that results from imperfections in the centering of the beam in the aperture.

Uncertainty in the measurement of relative aperture transmittance results from two primary sources: uncertainty in the centering of the beam in the apertures, and instability in the beam power and detector responsivity. The uncertainty in the aperture centering is determined experimentally, by physically

translating the external circular aperture in front of the transmittance measurement detector, and measuring the change in the response of the stationary detector. An upper bound is placed on the aperture misalignment; a 0.5 mm bound is currently used. The detector response when the aperture is well centered is compared with the response when the aperture is offset by the misalignment bound, the change in the response is then used as a bound for the centering uncertainty. Uncertainty that results from instability in the source power and detector is random, so is reduced by averaging multiple measurements. The main source of instability in the system is temperature fluctuations, which affect all of the elements of the calibration system to various degrees. Therefore, measurements with a lower uncertainty are obtained when the environment is more stable.

The detector is assumed to be perfectly uniform within its apparent aperture; variations in its spatial uniformity can affect the measured result, but the uncertainty caused by the detector's uniformity is assumed to be small compared to the uncertainty in the measurement. To minimize the uncertainty resulting from detector nonuniformity, the detector used for the circular aperture transmittance measurements have a spatial uniformity variation less than 1 % within their apparent aperture. When the test detector itself is used to measure its own aperture's transmittance, its uniformity can cause additional uncertainty in the measurement. In this case, the uncertainty will be at least partially accounted for when the end user adds the uncertainty due to the test detector's alignment, which is based in part on the spatial uniformity of the detector.

The result of imperfect detector uniformity can be better quantified once the actual beam profile is obtained with sufficient resolution; currently, an actual beam profile with sufficient dynamic range to resolve the details of the power distribution in the beam tails is not available. However, a theoretical example can help. Assume that 99.9 % of the beam's power is concentrated in the center of the detector, and 0.1 % is in a second beam that is produced by a multiple reflection. When a large aperture is placed in front of the detector, the second beam enters the detector. But when a small aperture is used, the second beam is blocked. The transmittance of the smaller aperture relative to the larger aperture is then $99.9 / 100 = 0.999$. Consider that because of nonuniformity in the detector, the second beam falls on an area where the responsivity is 1 % lower than that in the detector's center. The ratio would then be $99.9 / (99.9 + 0.1 \cdot 0.99) = 0.99901$, for a change of 0.001 %, which is negligible.

E.2 Relative Transmittance Measurement for Circular Apertures

The relative aperture transmittance for circular apertures is currently measured using a detector with an apparent aperture that is at least as large as the largest circular reference aperture (which has a 10 mm diameter), and has good spatial uniformity within its apparent aperture. The detector must also be sensitive to the beam's wavelength, and have low noise and good stability over a period of a few minutes. The circular apertures are placed in front of the large detector for the measurement; the detector itself is held stationary. If the detector has a built-in, circular, limiting aperture, then the built-in aperture can be used as the reference aperture, for the purpose of the transmittance measurements.

Using a detector with a single, flat, optically sensitive element that can be placed very close to (or in contact with) the circular apertures ensures that the detector is illuminated by all of the transmitted power, at the expense of its field of view. Placing the detector a distance behind the aperture can result in some of the transmitted power being diffracted away from the detector by the aperture's edge, so is less desirable. However, sometimes trap detectors are used even though their detecting element(s) must be some distance behind the aperture, because their superior spatial uniformity and reduced or eliminated back-reflection is more important than the small loss of diffracted power that may result; using trap detectors with apparent apertures that are larger than the reference aperture helps reduce any potential power loss. Since the resulting configuration can have a large field of view, a dark and unchanging

background around the beam is necessary. When a single-element detector or trap with a back reflection is used, it is placed in the beam at a small angle to prevent the back-reflection from interfering with the incident beam. Currently, a Si trap detector with a large apparent aperture is used with visible wavelengths, and a large single-element Ge detector is used at near-infrared wavelengths.

The relative transmittance of the aperture that is being measured is determined by dividing the amount of power transmitted with the measured aperture by the amount of power transmitted with the reference aperture, when both apertures are in turn illuminated by the same stable beam. The calculated transmittance is therefore relative to the transmittance of the reference aperture, and is appropriate only for the particular beam used in the measurement. The circular apertures are placed in the detector plane, where the detector's entrance apertures are also placed, so that the circular apertures are illuminated by the same profile beam that illuminates the entrance apertures of the detectors being calibrated. Since the measured circular apertures are all smaller than the reference aperture, their relative transmittance is always less than one. Therefore, a measurement result that returns a transmittance that is greater than one can be corrected to a value of 1. Apertures that are larger than the reference aperture can have a relative transmittance that is greater than 1.

The relative transmittance of several fixed, precision, circular apertures is measured with the system. The apertures are made of a thin ceramic material with precision circular holes. To prevent any optical power from penetrating the diffusely translucent ceramic material, the ceramic is painted with a black epoxy spray paint for opaqueness, and then further coated with a quality flat-black paint to reduce reflections. The precision apertures have diameters from 1 to 10 mm, in 1 mm steps. The 10 mm diameter aperture is usually used as the reference aperture.

The differences in the transmittance of the apertures is usually very small, especially for the larger apertures; therefore the measurement of relative transmittance must be performed very carefully. For example, the magnitude of the optical power drift is typically larger than the change in transmittance from one large aperture to another. However, the relative transmittance measurement is performed relatively quickly; a single measurement can be performed in less than a minute. And the drift in the optical power over such a time period is essentially linear, so linearly interpolating the power between bracketing reference aperture power measurements effectively reduces the effect of the power drift. To perform the interpolation, the period during which the power measurements occur is needed.

The transmittance of the smaller precision apertures relative to the reference aperture is determined with a series of measurements that are performed in such a way that the effects of drift in detector responsivity and beam power are minimized. The detector is illuminated throughout the measurement. First, the detector output P_1 is measured at time t_1 , using a no-baseline technique, with the reference aperture having diameter C in place (or with no additional aperture, if the detector's own circular limiting aperture with diameter C is used as the reference). Then the aperture under test with diameter D is placed on the detector, and the detector's response P_2 is measured at time t_2 . Then the test aperture is removed and the reference replaced, and the response P_3 is measured at time t_3 . Interpolating linearly in time, the relative transmittance is:

$$T_{RA}(D,C) = \frac{P_2}{P_1 + P_3} \left(\frac{t_3 - t_1}{t_2 - t_1} \right). \quad (\text{E-6})$$

The measurement of $T_{RA}(D,C)$ is repeated several times, and the mean and standard deviation are calculated. The mean estimates the relative transmittance, and the standard deviation of the mean, $u_M(D,C)$, is used in the uncertainty estimate. The measurements are performed rapidly to minimize the amount of drift encountered, so $P_1 \approx P_3$.

The measurements of the powers P_1 , P_2 , and P_3 contain uncertainties due to the imperfect centering of the apertures around the laser beam. The centering uncertainty is estimated in a separate experiment, described below, which provides an estimate of the standard uncertainty $u_p(X)$ due to imperfect centering in aperture X . In eq (E-6), the centering uncertainty in the measurement of P_2 is $u_p(D)$, and the uncertainty in P_1 and P_3 is $u_p(C)$. The uncertainties are uncorrelated, and approximating eq (E-6) with $T_{RA}(D,C) = P(D) / P(C)$, the standard uncertainty $u_{RA}(D,C)$ in $T_{RA}(D,C)$ is:

$$u_{RA}(D,C) = \sqrt{u_M^2(D,C) + \frac{u_p^2(D)}{P^2(C)} + \frac{P^2(D)}{P^4(C)} u_p^2(C)}. \quad (E-7)$$

The uncertainty due to the imperfect centering of the circular apertures around the beam is also measured. It is assumed that the apertures are centered around the beam with an error of less than 0.5 mm; experimentation has shown that the error in the aperture centering is less than this amount, since an error of 0.5 mm is clearly discernable by eye, even at infrared wavelengths (using an IR card). The measurement is performed to determine how much the aperture transmittance changes when the aperture is misaligned by this amount.

To measure the transmittance change due to the aperture centering error, the circular aperture with diameter d is placed in the detector plane, and centered around the beam as best as possible. The aperture is placed on a micrometer-driven mount so that it can be translated by a precise amount in the horizontal direction, and a detector is placed behind the aperture as in the primary transmittance measurement previously described. As before, the aperture is constantly illuminated with the stabilized beam. When the aperture is centered the detector's response is measured using a no-baseline technique and recorded as P_1 at time t_1 . Then the aperture is translated 0.5 mm in the positive horizontal direction, and the detector's response P_2 is measured at time t_2 . The aperture is then returned to center and the response P_3 is measured at time t_3 . The aperture is translated in the negative horizontal direction, and the detector's response P_4 is measured at time t_4 . Finally the aperture is again returned to the centered position, and the detector's response P_5 is recorded at time t_5 . Using linear interpolation, the relative transmittance change due to the positive translation, called $R_+(d)$, is given by

$$R_+(d) = \frac{P_4}{P_3 + P_5} \left(\frac{t_5 - t_3}{t_4 - t_3} \right). \quad (E-8)$$

The relative transmittance change for a negative translation, called $R_-(d)$, is calculated similarly:

$$R_-(d) = \frac{P_2}{P_1 + P_3} \left(\frac{t_3 - t_1}{t_2 - t_1} \right). \quad (E-9)$$

The measurements of $R_+(d)$ and $R_-(d)$ are repeated several times so that the mean and standard deviation can be calculated; the resulting means estimate the relative transmittance change, and the standard deviation of the means, called $u_{M+}(d)$ and $u_{M-}(d)$, are used to estimate the uncertainty.

In equation (E-7), the uncertainty u_p has units of power, so the relative measurements must be converted. To simplify the conversion, the small drift in the power with time is ignored, and the nominal power when the aperture is centered is called $P_0 = P_1 = P_3 = P_5$. Then the absolute change in the transmitted power due to the offset in the aperture's centering is estimated by $P_0(R_{\pm}(d) - 1)$, and the absolute standard deviation of the mean becomes $P_0 u_{M\pm}(d)$.

Because the actual centering uncertainty is assumed to be less than the ± 0.5 mm translation used in the measurement, the quantities $P_0(R_+(d) - 1)$ and $P_0(R_-(d) - 1)$ estimate the bounds of the change in transmitted power due to centering uncertainty. A rectangular probability distribution is assumed. The quantities can be positive or negative, because of the asymmetrical distribution of non-Gaussian power around the main Gaussian beam, such as that produced by scattering and multiple reflections, and error in the initial aperture centering. The resulting combined standard uncertainties due to the horizontal translations, called $u_{c+}(d)$ and $u_{c-}(d)$, are given by

$$u_{c+}(d) = P_0 \sqrt{u_{M+}^2(d) + \frac{(R_+(d) - 1)^2}{3}} \text{ and } u_{c-}(d) = P_0 \sqrt{u_{M-}^2(d) + \frac{(R_-(d) - 1)^2}{3}}. \quad (\text{E-10})$$

The greater of $u_{c+}(d)$ and $u_{c-}(d)$ is used as the standard uncertainty, $u_p(d)$, due to the imperfect centering of the circular aperture with diameter d around the beam. Using the greater uncertainty estimate helps compensate for any error in the initial centering of the aperture during the measurement of centering uncertainty.

Even when the transmittance measurement detector's built-in aperture is used to measure relative aperture transmittance, its aperture centering uncertainty must be calculated using an external aperture and a detector large enough to capture the entire transmitted beam, even when the aperture is displaced by 0.5 mm. If the detector with built-in aperture were used for the centering uncertainty measurement itself by translating the entire detector, the detector's spatial nonuniformity would affect the result when the detector is translated along with its aperture. Since the detector's spatial nonuniformity is likely to be much larger than the effect of the misaligned aperture, an external aperture that can be translated independently of the detector must be used.

E.3 Estimation of the Relative Aperture Transmittance for Detectors

When a detector has a true limiting aperture that is circular and is the same size as one of the measured circular apertures, the results of the measurements of circular aperture transmittance are used to determine the detector's transmittance relative to the standard. When a detector has a limiting aperture that is larger than the reference aperture, or when the detector has no true limiting aperture (such as that shown in Figure 6), the test detector's aperture transmittance is usually measured directly. Alternately, the test detector's transmittance can be estimated using bounding circular apertures. Bounding circular apertures are also used to estimate the transmittance of noncircular limiting apertures. Also, for the purpose of the calibration, an external circular aperture that fits within the detector's apparent aperture can be placed on the detector to form a true circular limiting aperture. However, the technique of adding an aperture to the detector is not usually used because it entails a modification of the test detector.

The transmittance of the test detector's aperture relative to the standard detector's aperture is used to correct the test detector calibration result. Three techniques, described below, are used to estimate test detector's relative transmittance. Each technique results in the estimated transmittance of the test detector's aperture, relative to the transmittance $T_{RA}(D,S)$ of the standard detector's aperture, with combined standard uncertainty $u_{RA}(D,S)$. In this context, D is the test detector's aperture, and S is the standard detector's aperture, but D and S do not necessarily represent the diameter of a circular aperture.

The standard detector's transmittance relative to the reference aperture's transmittance $T_{RA}(S,C)$, and its associated standard uncertainty $u_{RA}(S,C)$, are needed. These parameters are determined using a technique similar to one of the three described below, except instead of calculating the transmittance of the test

detector's aperture relative to the standard's aperture, the transmittance of the standard's aperture relative to the reference aperture is determined. For the LOCR's elliptical aperture, method 2 is used.

1. Estimation of the relative transmittance for detectors with a circular limiting aperture that is the same size as a measured circular aperture.

For these detectors, the measurement of the precision circular aperture's transmittance is used to estimate the transmittance of the detector's aperture, relative to the standard's aperture. Given that the detector has a circular limiting aperture with diameter D , and the reference is a circular aperture with diameter C , the relative transmittance of the detector's aperture $T_{RA}(D, C)$ is determined by performing the measurement described by eq (E-6). The corresponding combined standard uncertainty, $u_{RA}(D, C)$, is given by eq (E-7).

The standard detector's transmittance $T_{RA}(S, C)$ relative to the reference aperture's transmittance, and its associated standard uncertainty $u_{RA}(S, C)$, are previously determined. Then the transmittance $T_{RA}(D, S)$ of the test detector's aperture relative to the standard's aperture is given by

$$T_{RA}(D, S) = \frac{T_{RA}(D, C)}{T_{RA}(S, C)}. \quad (E-11)$$

The standard uncertainty in $T_{RA}(D, S)$ is given by

$$u_{RA}(D, S) = \sqrt{\frac{u_{RA}^2(D, C)}{T_{RA}^2(S, C)} + \frac{T_{RA}^2(D, C)}{T_{RA}^4(S, C)} u_{RA}^2(S, C)}. \quad (E-12)$$

2. Estimation of the relative transmittance for detectors with a noncircular or odd-sized circular limiting aperture that is bounded by measured circular apertures.

The noncircular or odd-sized circular aperture's relative transmittance is estimated using the measured transmittance of bounding, centered, circular apertures. Technique 3 is preferred for detectors that have an irregular apparent aperture but no limiting aperture, but this technique can be used with such detectors if necessary. The bounding apertures are the largest centered, circular aperture that is enclosed by the detector's apparent aperture, and the smallest centered, circular aperture that encloses the detector's apparent aperture. For example, the LOCR has a centered, approximately elliptical, limiting aperture with dimensions 8.5 by 9.5 mm; therefore the bounding circular apertures have diameters of 8 and 10 mm. The measured transmittance of the bounding circular apertures are combined to give an estimate for the transmittance of the noncircular aperture. Given that the bounding circular apertures have diameter D_1 and D_2 with mean relative transmittance $T_{RA}(D_1, C)$ and $T_{RA}(D_2, C)$ from eq (E-6); the estimated transmittance $T_{RA}(D, C)$ for the detector is given by the average (midpoint) of $T_{RA}(D_1, C)$ and $T_{RA}(D_2, C)$, or

$$T_{RA}(D, C) = (T_{RA}(D_1, C) + T_{RA}(D_2, C))/2. \quad (E-13)$$

The combined standard uncertainty $u_{RA}(D, C)$ of the transmittance estimate is determined by applying the law of propagation of uncertainty, and summing in quadrature with a term for the transmittance bounds. The resulting combined standard uncertainty is given by

$$u_{\text{RA}}(D, C) = \frac{1}{2} \sqrt{u_{\text{RA}}^2(D_1, C) + u_{\text{RA}}^2(D_2, C) + \frac{(T_{\text{RA}}(D_1, C) - T_{\text{RA}}(D_2, C))^2}{3}}. \quad (\text{E-14})$$

The transmittance of the detector's aperture relative to the standard's aperture can then be calculated using eq (E-11), with uncertainty calculated using eq (E-12).

3. Estimation of the relative transmittance for detectors with large apertures, and detectors that have no true limiting aperture.

The relative transmittance of detectors with large or unlimited apertures is measured directly, using a centered, circular aperture of an intermediate size. This technique is best for detectors that have no limiting aperture, because it includes the effects of the detector's partial response outside of its apparent aperture. However if the detector has a very slow or noisy response, method 2 can be used instead to obtain a more timely or less uncertain measurement. There is a small difference between the calculations for detectors with no limiting aperture, and those for detectors with a large aperture; the case of the detector with no limiting aperture, like that shown in Figure 6, is described first.

To measure the transmittance directly, the detector's response with no external aperture is compared to its response when the intermediate-sized external circular aperture is placed in front of the detector. The intermediate external circular aperture, which can be smaller than or equal to the reference aperture, must have a diameter equal to or smaller than the largest centered circular aperture that is completely enclosed in the detector's apparent aperture. The intermediate aperture is called I . To measure the detector aperture's transmittance relative to the intermediate aperture's transmittance, the detector is first placed with its entrance aperture in the detector plane, and is aligned to the beam in the same way that it will be aligned during absolute responsivity calibration. The mount with the intermediate aperture is then placed immediately in front of the detector, and is centered around the beam. The detector is then illuminated, and its transmittance $T_{\text{RA}}(D, I)$ relative to the intermediate aperture is measured as previously described for eq (E-6), except that aperture I replaces aperture C , and aperture D is the detector's unmodified aperture.

The combined standard uncertainty $u_{\text{RA}}(D, I)$ in the mean $T_{\text{RA}}(D, I)$ is the same as that given in eq (E-7), with a few changes to the uncertainty components. For $u_{\text{RA}}(D, I)$, the uncertainty $u_{\text{M}}(D, C)$ is replaced by $u_{\text{M}}(D, I)$, and the aperture centering term $u_{\text{C}}(C)$ is replaced by $u_{\text{C}}(I)$. To estimate the centering uncertainty term for the unmodified detector's aperture, $u_{\text{C}}(D)$, the smallest centered circular aperture that fully encloses the detector's apparent aperture, called D_2 , is used. However the centering uncertainty for the detector's aperture can be larger than $u_{\text{C}}(D_2)$, because of (a) the detector's partial response outside its apparent aperture, and (b) the asymmetrical power distribution outside the central Gaussian beam. To help compensate for the effect, the centering uncertainty for D_2 and all larger apertures are considered, and the largest centering uncertainty is used as $u_{\text{C}}(D)$ in eq (E-7). The centering uncertainty is generally smaller for larger apertures, but because of the asymmetrical power distribution is not necessarily so. While it is possible for the actual $u_{\text{C}}(D)$ to be larger, this estimation is considered adequate because the centering uncertainty for large apertures is usually insignificant compared to the uncertainty in the centering of the smaller intermediate aperture.

Estimation of the detector's unlimited aperture transmittance relative to the standard's aperture transmittance, called $T_{\text{RA}}(D, S)$, is complicated by the addition of the intermediate aperture. The relative transmittance is given by

$$T_{\text{RA}}(D, S) = \frac{T_{\text{RA}}(D, I) \cdot T_{\text{RA}}(I, C)}{T_{\text{RA}}(S, C)}. \quad (\text{E-15})$$

The combined standard uncertainty $u_{\text{RA}}(D,S)$ is given by

$$u_{\text{RA}}(D, S) = \sqrt{\frac{T_{\text{RA}}^2(I, C)}{T_{\text{RA}}^2(S, C)} u_{\text{RA}}^2(D, I) + \frac{T_{\text{RA}}^2(D, I)}{T_{\text{RA}}^2(S, C)} u_{\text{RA}}^2(I, C) + \frac{T_{\text{RA}}^2(D, I) T_{\text{RA}}^2(I, C)}{T_{\text{RA}}^4(S, C)} u_{\text{RA}}^2(S, C)}. \quad (\text{E-16})$$

When the detector's apparent aperture is larger than the reference aperture, the reference aperture itself can be used as the intermediate aperture, reducing the uncertainty in the measurement. Then, the relative transmittance is

$$T_{\text{RA}}(D, S) = \frac{T_{\text{RA}}(D, C)}{T_{\text{RA}}(S, C)}. \quad (\text{E-17})$$

The combined standard uncertainty $u_{\text{RA}}(D,S)$ is given by:

$$u_{\text{RA}}(D, S) = \sqrt{\frac{u_{\text{RA}}^2(D, C)}{T_{\text{RA}}^2(S, C)} + \frac{T_{\text{RA}}^2(D, C)}{T_{\text{RA}}^4(S, C)} u_{\text{RA}}^2(S, C)}. \quad (\text{E-18})$$

APPENDIX F. System DMM and LOCR Electronics: Calibration and Correction

All of the electrical systems that are directly involved in the detector calibration are traceable to the high-accuracy calibration system's digital multimeter (DMM), which is traceable to NIST standards. The system DMM itself is calibrated using an internal NIST calibration system, which returns the offset in the DMM's indicated readings. Since the offset in the DMM is reported but not nulled, a correction must be applied for every critical measurement that is performed with the DMM. This method of calibration adds significant complexity to electrical measurements performed in the laboratory. The correction is not applied at the time of the measurement, instead the indicating readings are recorded, and the correction is applied in post-processing where the calculations are fully documented. This technique also has the advantage that the acquisition system records readings that are consistent with those displayed on the device.

F.1 System DMM Calibration

The system DMM is calibrated using an internal NIST calibration system. The system measures the offset between the DMM's indicated voltage or resistance, and the actual voltage or resistance. The offset measurements are performed for each of the DMM's inputs and ranges, as requested. Only voltage and resistance measurements are needed so the DMM's current ranges are not calibrated. Generally, voltage measurements from 0 to ± 10 V using the 10 0.V scale are calibrated, and the offset for a resistance of 1,000 Ω is determined. The actual voltage V is the sum of the indicated voltage reading V_1 and voltage offset $V_c(V_1)$:

$$V = V_1 + V_c(V_1). \quad (\text{F-1})$$

The offset is measured for discrete voltages; offset values for voltages that fall between the calibrated values are linearly interpolated from the bracketing measured values. The bracketing calibrated voltages are V_1 and V_2 , where $V_1 < V_1 < V_2$. The interpolation is performed as follows:

$$V_c(V_1) = V_c(V_1) + \frac{V_1 - V_1}{V_2 - V_1} (V_c(V_2) - V_c(V_1)), \quad (\text{F-2})$$

where $V_c(V)$ is the calibrated offset voltage for indicated voltage V .

Similarly, the standard uncertainty $u_{V_c(V)}$ of V_c is also linearly interpolated from the standard uncertainties $u_c(V)$ of the bracketing values:

$$u_c(V_1) = u_c(V_1) + \frac{V_1 - V_1}{V_2 - V_1} (u_c(V_2) - u_c(V_1)). \quad (\text{F-3})$$

In addition to the uncertainty $u_c(V)$ of the calibration, the DMM experiences additional uncertainty with temperature and time. The difference between the operating temperature and the calibration temperature, and the interval since the calibration, add uncertainty to the DMM's measurements. The DMM's manufacturer has assessed and detailed these uncertainties in the DMM's operating manual [34]. The uncertainty is described in three parts: one relative to the DMM reading, one relative to the DMM range, and one relative to the reading that describes the traceability to NIST standards. The uncertainties are given as relative uncertainties in parts-per-million (ppm). The traceability to NIST electrical standards is included in $u_c(V)$, so only the first two uncertainties (ppm reading and ppm range) are relevant. The stated uncertainties are valid only for the operating mode and conditions specified in the document; the modes used in the high-accuracy calibration system conform to the specified requirements. The

uncertainty specifications include the device's temperature coefficient for a range of $\pm 5^\circ \text{C}$ about the calibration temperature. The laboratory temperature is normally within this range, so adding additional uncertainty is not usually needed. If the temperature is out of range, additional uncertainty is added to both the ppm range and ppm reading used in the example below.

Because of aging and drift in the DMM's internal components, the amount of uncertainty in the instrument's calibration increases with time, so the device's performance is specified for time intervals of 24 hours, 90 days, 1 year, and 2 years. The DMM is usually calibrated within 90 days of its use in the calibration laboratory, so the 90-day specification is usually used.

The uncertainties given by the manufacturer are an expanded uncertainty, but the coverage factor used is unclear. Therefore to be conservative, the specifications are assumed to be the limits of a rectangular interval, so are divided by the square-root of 3 to obtain the standard uncertainty. This is much more conservative than assuming a coverage factor of 2 or 3, and therefore dividing by 2 or 3 to get the standard uncertainty.

For example, when DC volts are measured using the DMM's 10 V range within 90 days of the previous calibration, the specified accuracy (uncertainty) is 4.1 ppm of reading plus 0.05 ppm of range, so the standard uncertainty u_{VS} is

$$u_{\text{VS}} = (10 \cdot 0.05 \cdot 10^{-6} + V_I \cdot 4.1 \cdot 10^{-6}) / \sqrt{3}. \quad (\text{F-4})$$

The combined standard uncertainty u_{V} of the DC voltage measurement is then the quadrature sum of the uncorrelated standard uncertainties u_{VS} and u_{VC} .

Table F.1 DMM voltage calibration values, for the front input terminals, and the 10 V range.

Indicated Voltage V_I (V)	Voltage Correction V_C (μV)	Standard Uncertainties		
		u_{VC} (μV)	u_{VS} (μV)	u_{V} (μV)
1.00	-2.70	3.0	2.7	4.0
2.00	-6.28	5.9	5.0	7.7
3.00	-9.30	8.9	7.4	11.5
4.00	-12.72	11.8	9.8	15.3
5.00	-16.55	15.0	12.1	19.3
6.00	-19.08	17.7	14.5	22.9
7.00	-22.40	20.7	16.9	26.7
8.00	-26.40	24.0	19.2	30.8
9.00	-29.25	27.0	21.6	34.6
10.00	-34.00	30.0	24.0	38.4
11.00	-35.86	33.0	26.3	42.2
11.76	-38.45	35.3	28.1	45.1

Table F.2 DMM resistance calibration values for the front input terminals and for the 1000 Ω range.

Indicated Resistance R_R (Ω)	Resistance Offset R_C (m Ω)	Standard Uncertainty			
		u_{RC} (m Ω)	u_{RS} (m Ω)	u_R (m Ω)	u_R (%)
1000	-174.8	4.3	2.7	5.0	0.00050

Example calibration results are shown for the DMM's front input terminals in Table F.1. Only positive voltages were calibrated in the example, but negative voltages are usually also calibrated. The uncertainty components u_{VC} and u_{VS} and combined standard uncertainty u_V are also given in the table. The 90-day DMM accuracy specifications are used to calculate u_{VS} .

In addition to the voltage calibrations shown on the table, the DMM's front terminal, four-wire resistance measurement was also calibrated at 1,000 Ω , as shown in Table F.2. The uncertainty in the resistance is assessed in the same way as the uncertainty in the voltage. Given in the table are the indicated resistance R_R , the resistance offset R_C , the standard uncertainty u_{RS} due to the NIST calibration, the DMM's specified uncertainty u_{RS} , and the combined standard uncertainty u_R of the resistance measurement. The combined standard uncertainty is given both in milliohms, and as a percentage of the indicated resistance. The specified uncertainty is the 90-day specification for the DMM's 1,000 Ω resistance range: 8 ppm of the reading, plus 0.5 ppm of the range. The DMM calibrations were performed at $T_C = 23 \pm 1^\circ\text{C}$.

F.2 Measurement of LOCR Electrical Power

The LOCR's internal electronics are calibrated by comparing its internally measured power to the power measured by the system DMM. Correction factors for the internal electronics are then stored in the LOCR. But this type of calibration is very time consuming, so is performed infrequently. However, the LOCR's calibration cannot be assumed constant with time and changes in the environment, so its calibration is checked for quality assurance whenever optical calibrations are performed, under the same conditions that exist during the optical calibration. The LOCR's calibration is only checked, it is not internally recalibrated to null the offset. Instead, the offset in the LOCR's internal calibration is measured. The correction for the LOCR's electrical calibration is combined with the correction for the DMM's measured voltage and resistance, to determine a single correction factor for the measured optical power.

During LOCR calibration, the system DMM is used to measure the power delivered to the LOCR receiver's heater. The electrical circuit consists of two resistors in series: the LOCR's heater resistor R_H , and the LOCR's standard resistor, R_S . Both resistors have a resistance of approximately 1,000 Ω . The precise resistance of the standard resistor, and the voltage drop around the heater and standard resistors, called V_H and V_S , respectively, are measured with the DMM. First, R_S is measured with no power applied to the circuit. Then the LOCR's electronics apply a current, and V_H and V_S are measured to determine the power P_H delivered to the heater. The heater power is given by

$$P_H = \frac{V_H V_S}{R_S}. \quad (\text{F-5})$$

The LOCR heater's resistance is within 1 % of the standard resistor's resistance; so to simplify the following uncertainty analysis, they are considered equal. Thus $R_H = R_S = R$, $V_H = V_S = V$, and the applied electrical power can be simplified to

$$P = \frac{V^2}{R}. \quad (\text{F-6})$$

The amount of absorbed optical power is determined by the change in electrical power that occurs when the optical power is applied, therefore the uncertainty in the difference between two measured electrical powers is needed. The LOCR's electronics maintain the receiver at a constant temperature by applying electrical power to the heater, so the change in applied electrical power necessary to maintain the constant temperature when optical power is applied, is equal to the amount of absorbed optical power. When operating with no absorbed optical power, the electrical system applies power P_1 to the receiver to maintain its temperature at the desired level. A nominal unilluminated power of 2 mW is used, so $P_1 = 2 \text{ mW}$. When optical power is absorbed, the electrical power must drop to P_2 to maintain the same temperature. The amount of optical power absorbed, called P_o , is therefore given by the following, where V_1 is the voltage applied when no optical power is present, and V_2 is the voltage applied when the optical power is present:

$$P_o = P_1 - P_2 = \frac{V_1^2 - V_2^2}{R}. \quad (\text{F-7})$$

The offset in the indicated voltages and resistance cause an offset in the power calculated from eq (F-4). The offset in the power can be determined by adding the absolute offsets dV_1 , dV_2 , and dR to V_1 , V_2 , and R , respectively. The optical power can then be expressed by

$$P_o(dV_1, dV_2, dR) = \frac{(V_1 + dV_1)^2 - (V_2 + dV_2)^2}{R + dR}. \quad (\text{F-8})$$

Since the DMM's readings have an offset, we desire the more accurate $P_o(V_{c1}, V_{c2}, R_c)$, where V_{c1} , V_{c2} , and R_c are interpolated from the calibration summarized in Tables F.1 and F.2, but $P_o(0,0,0)$ is what is actually measured. We desire a correction factor k_1 such that $P_o(V_{c1}, V_{c2}, R_c) = k_1 \cdot P_o(0,0,0)$. The correction factor is given by

$$k_1(V_{c1}, V_{c2}, R_c) = \frac{P_o(V_{c1}, V_{c2}, R_c)}{P_o(0,0,0)}. \quad (\text{F-9})$$

The uncertainty in k_1 is determined using the law of propagation of uncertainty. The uncertainty u_{k1} in k_1 , is given by

$$u_{k1} = \frac{1}{P_o(0,0,0)} \sqrt{\left(2 \frac{V_1 + V_{c1}}{R + R_c} u_{v1}\right)^2 + \left(2 \frac{V_2 + V_{c2}}{R + R_c} u_{v2}\right)^2 + \left[\left(\frac{(V_2 + V_{c2})^2 - (V_1 + V_{c1})^2}{(R + R_c)^2}\right) u_k\right]^2} + C. \quad (\text{F-10})$$

The term C represents the components that are correlated. When $V_1 \neq V_2$, the uncertainties are assumed to be uncorrelated, so $C = 0$. However, when $V_1 \approx V_2$, the uncertainties in the voltages are correlated, and omitting C will result in an inflated uncertainty estimate.

The covariance term in the law of propagation of uncertainty is estimated by $u(V_1, V_2) = k_{VC}u_{V1}u_{V2}$, where the term k_{VC} is the estimated voltage correlation. The other covariance terms are zero since the uncertainty in R is uncorrelated. In the limiting case where $V_1 = V_2$, the uncertainties are perfectly correlated so $V_{C1} = V_{C2}$, and $k_{VC} = 1$; when there is no correlation, $k_{VC} = 0$. Noting that $P_o(0,0,0)^2$ has been factored out of the denominator, C is then given by

$$C = -8 \frac{(V_1 + V_{C1})(V_2 + V_{C2})}{(R + R_c)^2} k_{VC} u_{V1} u_{V2}. \quad (F-11)$$

The correlation in the voltages is estimated from

$$k_{VC} = 1 - 2 \frac{|V_{C1} - V_{C2}|}{|V_{C1} + V_{C2}|}. \quad (F-12)$$

When $k_{VC} < 0$, the voltages are considered uncorrelated, and $k_{VC} = 0$ is used instead.

The correction factor and uncertainty, calculated using the example data given in Tables F.1 and F.2, are shown in Table F.3. Uncertainty estimates are given for uncorrelated voltages, 100 % correlated voltages, and using the estimated voltage correlation. In practice, the result for the estimated voltage correlation is used, but the other results are given to show by how much the uncertainty changes for the different degrees of correlation.

F.3 Calibration of LOCR Power

The LOCR's internal electrical heater power measurements are calibrated using the "ACRCAL.EXE" program. This program calibrates the LOCR's internal electronics over their entire operating range, so is very time consuming, and environmental drift may occur during the lengthy calibration. To better assess the LOCR's uncertainty at the time of an actual optical calibration, the uncertainty of its internal electronics is assessed using the calibration verification program "TCALVER.EXE". The calibration is verified at only the power levels of interest in the detector calibration, and in an environment consistent with that present during the detector calibrations. The verification program returns the offset in the internal electronic's power measurement, based on measurements performed with the system DMM, as described in the previous section. The offsets are used to generate a correction factor k_2 for the LOCR's internal electronics and its combined standard uncertainty u_{k2} . The measurement is performed several

Table F.3. Examples of the correction factor k_1 and its uncertainty, for several optical powers.

Optical Power P_o (mW)	Electrical Power P_M (mW)	Voltages			Cor- relation k_{VC}	Result k_1	Standard Uncertainty $u_{k1} \times 10^6$		
		V_M (V)	V_{CM} (μ V)	u_{VM} (μ V)			Uncor- related	100 % Cor- related	Estimated Cor- relation
All	2.00	1.414	-4.183	5.533	N/A	N/A	N/A	N/A	N/A
1.00	1.00	1.000	-2.700	4.000	0.657001	1.000168	18	9	13
0.25	1.75	1.323	-3.856	5.195	0.933503	1.000168	84	9	23
0.10	1.90	1.378	-4.055	5.400	0.974212	1.000168	215	9	36
0.01	1.99	1.411	-4.170	5.519	0.997876	1.000168	2,605	10	120

N/A - Not applicable

times to assess the random component of the uncertainty, but the uncertainty in the electrical measurements themselves is accounted for elsewhere, in k_1 , so the indicated electrical powers are assumed to be accurate here.

When the LOCR is used to measure a single optical power, the offset of its internal electronics at two power levels, P_1 and P_2 , are needed. Since P_1 is always the same (2 mW), only one additional offset needs to be measured for each additional power level. However, the actual absolute power levels can fluctuate for several reasons, including drift in the environment and the absolute optical power, so the LOCR's calibration is assessed for a region of power levels that includes the potential drift. The drift in the nominal absolute powers is limited to an amount d_p ; if either power level P_1 or P_2 change from their nominal values by more than d_p , the LOCR and/or optical source is readjusted. The power drift d_p can therefore be considered as the bound of a rectangular distribution around the nominal power P . A d_p of 1 % is easily achieved; at higher powers, a d_p of 0.1 % can be achieved.

For each nominal power P , the LOCR calibration offset is measured at several points within the interval $P \pm d_p$, and a statistical analysis is performed. To complete the calibration in a reasonable amount of time, the power offset is usually measured at only three levels, P , $P + d_p$, and $P - d_p$. The offset at each level is measured at least once, and the mean offset and standard deviation of the mean are calculated for all the measurements within the interval. The mean offset $P_c(P)$ for indicated power P is such that the corrected power is $P + P_c(P)$. The standard deviation of the offsets is the standard uncertainty $u_{pc}(P)$ of the offset; the standard deviation of mean is not used because we desire the variation that occurs with the variation in the nominal power, not the uncertainty of the nominal power measurement. The results for a typical LOCR calibration at several power levels are shown in Table F.4; a d_p of 1 % was used. In the

Table F.4 Typical LOCR power calibration check results.

Nominal Power P (mW)	Indicated Power (mW)	Measured Offset (nW)	Average Power Offset $P_c(P)$ (nW)	Standard Uncertainty $u_{pc}(P)$ (nW)
2.00	2.02	-14.1	-14.7	1.0
	2.00	-14.0		
	1.98	-15.8		
1.90	1.92	-13.4	-14.6	2.1
	1.90	-13.3		
	1.88	-16.9		
1.75	1.77	-17.7	-14.0	3.6
	1.75	-10.5		
	1.73	-13.9		
1.00	1.01	-10.1	-11.3	1.1
	1.00	-12.0		
	0.99	-11.9		

table, a single measurement at each indicated power is shown, but in practice the measurements should be performed multiple times. For simplicity, such multiple measurements are not pre-averaged; but if it becomes desirable to reduce this uncertainty, pre-averaging could be employed.

In optical power calibrations, the change in indicated power is calculated. In this case, the optical power, P_o , is expressed as the following function:

$$P_o(P_1, P_2) = P_1 - P_2. \quad (\text{F-13})$$

To correct for the offset in the indicated optical power, the indicated electrical powers in eq (F-13) are replaced by the indicated powers P plus the power corrections $P_c(P)$, from Table F.4. We desire the multiplicative correction factor k_2 that converts the indicated optical power to the corrected optical power, so

$$k_2 = \frac{P_o(P_1 + P_c(P_1), P_2 + P_c(P_1))}{P_o(P_1, P_2)} = 1 + \frac{P_c(P_1) - P_c(P_2)}{P_1 - P_2}. \quad (\text{F-14})$$

Using the law of propagation of uncertainty, the combined standard uncertainty u_{k2} of k_2 is given by

$$u_{k2} = \frac{\sqrt{u_{PC}(P_1)^2 + u_{PC}(P_2)^2 + C}}{P_1 - P_2}. \quad (\text{F-15})$$

As before, C represents the components that are correlated. The covariance term is estimated by $u_{PC}(P_1, P_2) = k_{PC}u_{PC}(P_1)u_{PC}(P_2)$, where k_{PC} is the estimated correlation in the powers. In the limiting case where $P_1 = P_2$ the uncertainty is perfectly correlated, so $k_{PC} = 1$; and when the uncertainty is uncorrelated, $k_{PC} = 0$. The correlation term C is then given by

$$C = -2k_{PC}u_{PC}(P_1)u_{PC}(P_2). \quad (\text{F-16})$$

The correlation in the powers is estimated by

$$k_{PC} = 1 - 2 \left| \frac{P_1 - P_2}{P_1 + P_2} \right|. \quad (\text{F-17})$$

When $k_{PC} < 0$, the powers are considered uncorrelated and $k_{PC} = 0$ is used instead.

Table F.5 Example LOCR power calibration correction factor and uncertainty.

Optical Power (mW)	k_{PC}	k_2	Standard Uncertainty $u_{k2} \cdot 10^6$		
			Uncorrelated	Correlated	Estimated
1.00	0.738426	0.999997	1	0	1
0.25	0.951220	0.999997	15	10	11
0.10	0.993174	0.999999	23	11	11
0.01	0.999319	0.999999	149	11	12

The correction factor k_2 and its uncertainty are given in Table F.5 using the data from Table F.4. The values for a nominal power of 1.99 mW, used in the 0.01 mW calculations, are interpolated from the results for powers of 1.9 and 2.0 mW. Three uncertainty estimates are given; the uncorrelated result assumes no correlation in the powers, the correlated result assumes 100 % correlation, and the estimated correlation uses the estimated value of k_{PC} . In practice the estimated correlation result is used, but the other results are given to show by how much the uncertainty changes for the different degrees of correlation.

F.4 Combined LOCR Power Measurement Correction

The combined correction factor for the LOCR's power measurements is the superposition of the DVM voltage and resistance measurement corrections, and the LOCR calibration correction. Therefore the combined correction factor k_L is given by

$$k_L = k_1 k_2. \quad (F-18)$$

The combined standard uncertainty u_{kL} of k_L is estimated using the law of propagation of uncertainty, assuming that the uncertainties are uncorrelated:

$$u_{kL} = \sqrt{k_2^2 u_{k1}^2 + k_1^2 u_{k2}^2}. \quad (F-19)$$

The correction factor and its standard uncertainty are given in Table F.6. Two uncertainty estimates are given; the uncorrelated result assumes that the uncertainties are uncorrelated, the correlated result assumes that the uncertainty in the voltages and powers are perfectly correlated, and the estimated result uses the estimated correlation. The estimated result is used in practice, and the other results are given to show by how the uncertainty changes for different degrees of correlation. The 0.01 mW example shows that the LOCR electronics can accurately measure such low powers, but random uncertainty in the optical power measurement itself would likely be prohibitively large.

Table F.6 Combined LOCR electrical calibration correction factor and uncertainty.

Optical Power (mW)	k_L	Standard Uncertainty u_{kL} (%)		
		Uncorrelated	Correlated	Estimated
1.00	1.000165	0.0018	0.0009	0.0013
0.25	1.000165	0.0085	0.0013	0.0025
0.10	1.000167	0.0216	0.0014	0.0038
0.01	1.000167	0.2609	0.0015	0.0121

APPENDIX G. LOCR Optical Receiver Alignment: Uncertainty

It is usually assumed that with cryogenic radiometers, the optical receiver's absorptance is uniform within its apparent aperture. However, a recent publication [8] suggests that difficulties encountered in measurement of optical receiver absorptance may be due, at least in part, to spatial nonuniformity in the receiver's absorptance. Subsequent experiments performed on the LOCR confirmed that there is significant non-uniformity in its optical receiver's absorptance. This nonuniformity, combined with the imperfect centering of the laser beam within the receiver, adds additional uncertainty to power measurements made with the device. In the high-accuracy calibration system, this uncertainty is accounted for in the LOCR alignment uncertainty.

G.1 Philosophy of Receiver Uniformity Measurement

Unlike the measurements performed for the external publication, the LOCR measurements were performed with the LOCR in normal operating mode, with its optical receiver at normal operating temperature and pressure. Since a greater nonuniformity was encountered in the LOCR measurements than was inferred in the publication, the nonuniformity may increase when the optical receiver is at cryogenic temperature. The source of the nonuniformity in absorptance may be due to spatial variations in the paint thickness, adhesion, consistency, or contamination. It may also be due to cracks in the paint that open wider when the receiver is at cryogenic temperature. These explanations are speculation, but no matter what the source of the nonuniformity, the measurements show that a significant nonuniformity is present. Since the magnitude of the nonuniformity can be affected by the size and wavelength of the laser beam used, the LOCR's uniformity is measured after absolute power calibration are performed, using the same laser beam that was used in the calibrations. When a sufficient history of uniformity measurements is acquired, it may be possible to eliminate the measurement and use historical values. However, performing the measurement provides another level of quality assurance, since the receiver uniformity may be degraded at any time.

The LOCR optical receiver's spatial uniformity is estimated by measuring relative power at several different locations around the center of its optical receiver, while the LOCR is in normal operating mode. Since the optical power and LOCR's thermal sensitivity are constant, any variation in the measured responsivity must be due to variation in the amount of power absorbed by the receiver. The amount of power absorbed can vary for several reasons: nonuniformity in the LOCR's window transmittance, a change in the aperture transmittance as the beam location within the aperture is varied, or nonuniformity in the receiver's absorptance. When the first two sources of variation are removed from the measured variation, any remaining variation must therefore be due to nonuniformity in the receiver's absorptance.

The nonuniformity in receiver absorptance results in uncertainty in absolute power measurements because the location of the beam in the LOCR's receiver (the LOCR's alignment) is itself uncertain. The location of the beam during the absorptance measurement is unknown, but the location was most likely not exactly the same location used in the absolute power measurements. A conservative alignment uncertainty of 1 mm is used, because such a large error in the centering of the beam in the receiver would be obvious, and therefore should not occur in practice. We are therefore concerned with the variation in the optical receiver's absorptance within a centered, circular area with a radius of 1 mm. This radius could be reduced to perhaps 0.5 mm in the future, if new receiver absorptance measurements are performed with a known uncertainty in centering.

Because of limitations in the current calibration system, and the necessity of calibrating the detector in a timely manner, the variation in responsivity within the entire circular area is not completely measured. Instead, a statistical sampling of the relative absorptance within the area is used to estimate the variation

within the entire circular region. A vertical translation of the beam's location in the receiver cannot be performed. To translate the beam vertically, the entire optical source breadboard must be raised or lowered. Currently, this involves manually adjusting the five leveling posts that support the optical source breadboard. Performing the manual adjustment requires a significant amount of time, and the beam's location in the detector plane swings wildly as the posts are adjusted one at a time. The time required to adjust the posts, in addition to the time required for the LOCR to stabilize after the wild swing in beam location, is prohibitive because the drift in the absolute power during this time is too large to be effectively compensated for by linear interpolation. Therefore, vertical translations cannot currently be used in the measurement of receiver uniformity. In the future, if the manual leveling posts are replaced with motion-controlled jack stands, the vertical translation could be performed quickly, and the receiver's uniformity more completely measured.

G.2 Receiver Uniformity Measurement: Technique and Analysis

The amount of nonuniformity is small compared to the amount of drift in the optical power, so is difficult to measure. To minimize the affect of the drift in the absolute power, a relative measurement is performed quickly, typically in less than one minute. As described in Appendix E, the drift in the optical power over such a time period is essentially linear; linearly interpolating the power between bracketing measurements effectively reduces the effect of the power drift. To perform the interpolation, the time that the power measurements occur is also needed.

The LOCR's responsivity relative to its responsivity when the beam is located at the receiver's nominal center is measured. The beam is first located at the nominal center, and the LOCR's response P_1 is measured at time t_1 . Then, the beam is translated to location (x,y) on the receiver, and the response P_2 is measured at time t_2 . Finally, the beam is translated back to the nominal center, and the response P_3 is measured at time t_3 . By interpolating the nominal center response in time, the responsivity at location (x,y) , relative to the response at the nominal center, is given by $r(x,y)$, where

$$r(x,y) = \frac{P_2}{P_1 + P_3} \left(\frac{t_3 - t_1}{t_2 - t_1} \right). \quad (\text{G-1})$$

The responsivity $r(x,y)$ is measured several times, and the average of $r(x,y)$, $R(x,y)$, is used as the final responsivity at location (x,y) . The uncertainty $u_{R(x,y)}$ in the relative responsivity is dominated by the measurement reproducibility, so is given by the standard deviation of the mean of $R(x,y)$. If the measurement reproducibility is reduced in the future, assessment of the uncertainty in the power measurements themselves may be needed; but the uncertainty in the powers are strongly correlated, so the resulting uncertainty contribution would be small, as described in Appendix F.

Since a vertical translation in the beam's location cannot be performed at this time, samples taken at horizontal translations only are used. Currently measurements acquired at five locations, spaced 0.5 mm apart on a horizontal line through the receiver's nominal center are used. The locations are $x = -1, -0.5, 0, +0.5$, and $+1$ mm; the nominal center is located at $x = 0$ mm. The y -locations are always zero, so in the following description, the second subscript is dropped and $R(x) = R(x,y) = R(x,0)$.

To obtain the receiver absorptance, spatial variation in the aperture and window transmittance must be removed from $R(x)$. Variations caused by the aperture translation are measured to assess the uncertainty in the aperture transmittance measurement, described in Appendix E; spatial variation in the window transmittance is measured during window transmittance measurement, described in Appendix H. The corresponding window transmittance values are called $T_w(x)$, with standard uncertainty $u_{T_w}(x)$, and the

aperture transmittance values are called $T_A(x)$, with standard uncertainty $u_{TA}(x)$. Like the responsivity measurements, the transmittances are also measured relative to the transmittance at the nominal center, so each measurement has a value of 1 at $x = 0$. The receiver absorptance $A_R(x)$ is then given by

$$A_R(x) = \frac{R(x)}{T_W(x)T_A(x)}. \quad (G-2)$$

The combined standard uncertainty $u_{AR}(x)$ of the relative receiver absorptance measurement at location x is given by

$$u_{AR}(x) = \sqrt{\frac{u_R^2(x)}{T_W^2(x)T_A^2(x)} + \frac{R^2(x)u_{TW}^2(x)}{T_W^4(x)T_A^2(x)} + \frac{R^2(x)u_{TA}^2(x)}{T_W^2(x)T_A^4(x)}}. \quad (G-3)$$

The terms $T_W(x)$ and $T_A(x)$ are measured during the final measurement of window transmittance, and the measurements of relative aperture transmittance, respectively. The relative aperture transmittance $T_A(x)$ is measured in the same way as the relative aperture transmittance's aperture centering uncertainty, described in Appendix E, but two additional measurements are required. For the relative aperture transmittance measurement, the relative transmittance for only an aperture displacement of ± 0.5 mm is measured; for the receiver uniformity measurement, additional measurements at a displacement of ± 1.0 mm are also needed. Since the LOCR's receiver has an elliptical apparent aperture with major and minor axes of approximately 8.5 by 9.5 mm, the results for a 9 mm diameter circular aperture are used.

Similarly, additional window transmittance measurements similar to those described in Appendix G are needed. For this purpose, the change in window transmittance relative to the transmittance at the window's nominal center is needed, not the absolute transmittance measured for the window transmittance calculation. Instead of measuring the absolute transmittance, by moving the window in and out of the beam, the transmittance at the displaced location is measured relative to the transmittance when the window is nominally centered, by displacing the window and returning it to the nominal center location. An analysis analogous to that described in Appendix E for the aperture centering uncertainty is used. This measurement is performed only during the final window transmittance measurement, using the same nominal window center location that was used in the calibrations. Both the aperture displacement and the window displacement are measured several times, and the means of the measurements are used as $T_A(x)$ and $T_W(x)$. The standard deviation of the means are the standard uncertainties $u_{TA}(x)$ and $u_{TW}(x)$.

The uncertainty in LOCR alignment was discovered recently, so the additional measurements described above were not performed during earlier calibrations. Therefore only the receiver uniformity results from later calibrations, performed at visible wavelengths, are available. In all future test detector calibrations, the aperture offset and relative window transmittance measurements described above will be performed.

Since the locations of the nominal center used in the calibration runs and receiver absorptance measurement are unknown, a correction for the LOCR's alignment cannot be applied. Only the combined standard uncertainty due to the LOCR's alignment, called u_{LA} , is needed. Using a "between and within" analysis, two components of uncertainty, called u_{LA1} and u_{LA2} , contribute to u_{LA} . The standard uncertainty u_{LA1} is the standard deviation of the 5 values of $R(x)$. The standard uncertainty u_{LA2} is the average of the uncertainties $u_R(x)$, for $x = -1.0, -0.5, 0.5$, and 1.0 (the uncertainty $u_R(0)$ is omitted because the

Table G.1 Example LOCR alignment uncertainty data.

	$x = -1.0$ mm		$x = -0.5$ mm		$x = 0^*$	$x = 0.5$ mm		$x = 1.0$ mm	
	Value	Std. Unc.	Value	Std. Unc.	Value	Value	Std. Unc.	Value	Std. Unc.
$\lambda = 633$ nm									
$R(x)$	0.999900	0.000020	0.999902	0.000025	1.000000	1.000181	0.000023	1.000226	0.000023
$T_W(x)$	0.999974	0.000013	0.999987	0.000012	1.000000	1.000021	0.000012	1.000043	0.000011
$T_A(x)$	0.999980	0.000009	0.999986	0.000006	1.000000	1.000002	0.000006	1.000008	0.000009
$A_R(x)$	0.999946	0.000025	0.999929	0.000028	1.000000	1.000157	0.000027	1.000176	0.000027
$\lambda = 515$ nm									
$R(x)$	0.999780	0.000001	0.999924	0.000015	1.000000	1.000033	0.000005	0.999975	0.000008
$T_W(x)$	0.999924	0.000012	0.999962	0.000011	1.000000	0.999982	0.000011	0.999964	0.000010
$T_A(x)$	0.999984	0.000001	0.999996	0.000001	1.000000	1.000001	0.000002	0.999995	0.000001
$A_R(x)$	0.999872	0.000012	0.999965	0.000019	1.000000	1.000050	0.000012	1.000016	0.000013
$\lambda = 458$ nm									
$R(x)$	1.000302	0.000033	1.000365	0.000034	1.000000	0.999874	0.000047	0.999850	0.000020
$T_W(x)$	1.000038	0.000018	1.000019	0.000011	1.000000	1.000048	0.000011	1.000095	0.000009
$T_A(x)$	1.000003	0.000001	1.000003	0.000001	1.000000	1.000000	0.000001	1.000000	0.000001
$A_R(x)$	1.000261	0.000038	1.000343	0.000036	1.000000	0.999826	0.000048	0.999754	0.000022

* - Values are measured relative to the value at $x = 0$, so there is no uncertainty at $x = 0$.

uncertainty at location $x = 0$ is unknown). The combined standard uncertainty u_{LA} is then given by the quadrature sum of u_{LA2} and u_{LA2} .

Example calculations, using the data from Table G.1, are shown in Table G.2. Measured values of u_{LA} have varied from approximately 0.007 to 0.03 % at different wavelengths; the uncertainty is therefore very significant, and can even dominate the calibration uncertainty at some wavelengths. These values are for a beam diameter of 2 mm beam. The combined standard uncertainty u_{LA} is likely to increase for smaller diameter beams, and decrease for larger beams. Environmental instability contributed to the size of the uncertainty, especially at 458 nm.

Table G.2 Example LOCR alignment uncertainty calculation.

λ (nm)	u_{LA1}	u_{LA2}	u_{LA}
633	0.000117	0.000027	0.000120
515	0.000068	0.000014	0.000069
458	0.000260	0.000036	0.000262

APPENDIX H. Window Transmittance: Estimation and Uncertainty

The transmittance of the Brewster's angle window is the ratio of the amount of optical power that is incident on the cryogenic radiometer's optical receiver when the radiometer's window is in place, relative to the amount of power incident on the receiver when the window is removed. A theoretical transmittance for the LOCR Brewster's angle window can be calculated from the wavelength of the laser light and the properties of the window's glass, but as with the relative aperture transmittance, imperfections in the real system cause the theoretical calculation to be in error by an amount that is significant in the high-accuracy calibration system. Therefore, like the relative aperture transmittance, the window transmittance must be measured. The window transmittance is measured each time the LOCR is used for a calibration because the transmittance depends on many factors, which include the window's type, surface quality, alignment, orientation, contamination, scatter centers, and other imperfections; along with the specifics of the laser beam, such as its wavelength, profile, and radius of curvature.

H.1 Window Transmittance Measurement: Philosophy and Technique

The transmittance of the Brewster's angle window is measured using a test detector with little or no back-reflection and which has an apparent aperture that is at least as large as the LOCR's. The window assembly is mounted at the same location relative to the optical source as when it is on the LOCR, and the test detector is mounted at the same relative location as the LOCR's optical receiver. The test detector is fitted with two external aperture stops that simulate the apertures present in the LOCR; one aperture stop simulates the LOCR's limiting aperture, and the second simulates the first aperture in the LOCR's cryogenic shield. The test detector therefore has the same field of view as the LOCR's receiver. The test detector and aperture stops are mounted on a jig that contains a jack stand that allows adjustment of the assembly's height, without other translations. The window assembly is mounted on the transverse rail, which allows the window to be moved in front of the test detector's jig without other translations. The window is adjusted to Brewster's angle by minimizing the intensity of the reflection from the front-side of the window; the reflection from the back of the window is not used.

To determine the transmittance of the window, the power incident on the test detector when the window is in place is measured and divided by the power incident on the test detector when the window is removed. Since there is a vertical (and slight angular) offset of the transmitted beam when the window is in place, the jig's jack stand is adjusted to keep the beam in the center of the apertures. The jig also contains a translation stage for lateral adjustment, if necessary. To reduce the effect of drift in the applied power and detector responsivity, the measurement is performed quickly, and the power present at the time of the measurement with the window in place is interpolated from bracketing measurements made without the window, as previously described in relation to the power calibration and relative aperture transmittance measurements. The transmittance measurement is repeated several times, and a mean and standard deviation of mean are calculated.

Since there is always some loss in the window, the transmittance should be less than one. Therefore, a measured transmittance that is greater than one indicates a problem with the measurement; using a detector with a strong back reflection can cause such a problem. When using a three-element reflection trap as the test detector, the NIST window mount's transmittance was frequently greater than one until the inside of the mount was painted black. Using a transmission trap in the measurement helps ensure that any remaining problem with back reflections is eliminated.

The uncertainty in the window transmittance measurement is dominated by three factors: drift in the transmittance during the calibration, uncertainty in where the beam is located on the window, and noise in the transmittance measurement itself. These uncertainties are measured as described in the next section.

Other potential sources of uncertainty are insignificant compared to these factors, so are not included in the uncertainty analysis. Two such insignificant uncertainties are: imperfection in the scattered optical power measurement, and the change in transmittance due to the vacuum in the cryogenic radiometer; these factors are discussed first.

Calibration laboratories usually use one of two philosophies to account for the optical power that is scattered by the window. The technique used in the high-accuracy calibration system accounts for the scattered power with the transmittance measurement itself. Other laboratories directly measure the amount of scattered power, and generate a correction factor and uncertainty for the scattered power that is applied to the transmittance measurement. Measuring the scattered power requires additional sensors located inside the cryogenic radiometer, so devices without the additional sensors, such as the LOCR, cannot use the method.

With the LOCR, optical power that is scattered by the window is accounted for in the transmission measurement itself, by duplicating the field of view of the LOCR's optical receiver with the test detector that is used in the window transmission measurement [3]. Therefore, any scattered power that misses the LOCR's optical receiver also misses the detector used in the transmittance measurement. The loss due to scattering is therefore measured along with the window transmittance itself. Any difference between the two scattering losses, such as the additional scattering caused by the air, which is present in the transmittance measurement but not when the window is on the LOCR, is negligible when compared to the primary uncertainties in the measurement. One remaining source of error is scattered power that is reflected to the optical receiver by the walls of the connecting steel vacuum lines, which are present when the window mount is on the LOCR, but are absent during the transmittance measurement. Painting the inside of these tubes with a vacuum-compatible flat-black paint, or duplicating the tubing in the transmittance measurement would eliminate this concern; currently any error resulting from the potential reflection is assumed to be negligible.

Another potential source of uncertainty in the window transmittance measurement arises from the fact that the transmittance is measured with the window in air, but when on the LOCR, the back-side of the window is exposed to vacuum. The reflection from the second surface of the window is therefore slightly different when the window is on and off the LOCR. The theoretical change in the transmittance can be determined using the Fresnel equations to estimate the window's transmittance to light with polarization angle that is parallel to the plane of incidence [33]. The following results were produced using a quartz-glass index of 1.457, and a wavelength of 633 nm. To obtain a worst-case result, an air index of 1.0003, which is appropriate at sea-level but is slightly higher than the actual index present in the NIST-Boulder laboratory, is used in the calculations. Assuming that the window's front and back surfaces are perfectly parallel, when the window is in air and the light's incidence angle is at Brewster's angle, there is no reflection from either the front or the back of the window (for the parallel polarization). When the back of the window is exposed to vacuum, the transmittance changes by approximately $-3 \cdot 10^{-6} \%$. The situation is somewhat worse for the actual window, which has a small wedge angle. The windows currently used with the LOCR have a wedge angle of no greater than 5 minutes of arc (about 1.5 mrad). Assuming the worst-case wedge angle orientation, when the front-side of the window is at Brewster's angle, applying vacuum to the back changes the transmittance by about $\pm 6 \cdot 10^{-5} \%$, which is negligible when compared to the other uncertainty terms, which have a typical magnitude of $10^{-3} \%$, so the effect is ignored. The presence of the wedge also reduces the overall transmittance slightly, by approximately $3 \cdot 10^{-4} \%$ for the worst-case scenario; depending on the orientation of the wedge angle relative to the polarization angle, the loss can be smaller. The window's nonzero wedge angle can therefore reduce the maximum theoretical window transmittance, but any reduction is accounted for by the transmittance measurement. The presence of a surface contaminant on the glass can have a much greater effect, as described later.

The presence of the wedge and the vacuum on the back of the window also cause the direction of the transmitted beam to deviate slightly from the direction of the incident beam. For the perfectly parallel window, the vacuum on the back of the window causes the transmitted beam to be deviated by about 0.4 mrad relative to the incident beam; the worst-case deviation for the wedged window is about 4 mrad. But any resulting deviation is compensated for during the LOCR and test detector alignment, so does not significantly affect the accuracy of the measurement. The thickness of the window also causes a vertical translation of the transmitted beam, as described in the LOCR alignment section. The potential for the intensity of the reflection from the back of the window to change when the window is removed from the LOCR's vacuum is the reason why only the reflection from the window's first surface is used to adjust Brewster's angle; the intensity of the reflection from the front-side of the window is constant when the window is both on and off the LOCR. Stress-induced birefringence in the window glass would complicate this analysis, but the thick quartz window and low-stress window mount used are assumed to minimize any such birefringence to negligible proportions.

H.2 Transmittance Measurement: Uncertainty Analysis

The window transmittance and uncertainty used in the detector calibration are derived from multiple measurements of window transmittance performed both before and after the calibration runs, using the analysis described below. This analysis is the result of a trade-off between performing timely calibrations, and performing extremely accurate measurements of window transmittance. Given that some other laboratories measure the window transmittance at only a single location on the window, either before or after the calibrations are performed, this analysis is conservative in comparison. The uncertainty in the window transmittance measurement is dominated by three factors: drift in the transmittance during the calibration, uncertainty in where the beam is located on the window, and noise in the transmittance measurement itself. In the high-accuracy calibration system, all of these factors are included in the uncertainty analysis of window transmittance.

The transmittance drift during the calibration is frequently the largest uncertainty component. The transmittance measured after the calibration is almost always higher than that measured before the calibrations. The drift is likely due to an absorbate, such as water vapor, accumulating on the back-side of the window, when the window is not in use. When the window is then used on the LOCR, the absorbate is slowly released into the LOCR's internal vacuum, and the window's transmittance increases. To account for the affect, the window's transmittance is measured both before and after the calibrations are performed. To help reduce the magnitude of the drift, the window is placed on the LOCR and exposed to vacuum for some time before the initial transmittance measurement is performed. If the LOCR is used more often in the future, the back of the window will be exposed to the atmosphere for less time, and the drift should be reduced.

The window transmittance measurement is complicated by the fact that the transmittance is not uniform over the surface of the window. Also, there is uncertainty in the exact location where the laser beam strikes the window, as described in section 3.2.4. To account for the uncertainty in the beam's location, a matrix of transmittance measurements around the nominal beam location is acquired. A nine-point square matrix with a 1 mm spacing is used. The 1 mm spacing is used because the shift in the beam's location on the window, due to imperfections in the mounting hardware, is typically 1 mm or less. However, the beam diameter at the window can be smaller than 1 mm, so the matrix can be considered a statistical sampling of the transmittance, rather than a complete mapping of the transmittance in the region.

The result of the two transmittance measurements is two arrays of nine transmittance means, called $T_w(x,y)$, with their associated standard deviation of mean, called $u_w(x,y)$. As described in section 3.2.4, the nominal beam location in the first set of transmittance measurements can be different from the

nominal beam location in the second set of measurements, because of imperfections in the mounting hardware. Therefore, the transmittance value at a location (x,y) in the two matrices may not represent the same physical location on the window. The location shift would have to be reconciled if a more complex analysis is used, such as the center-weighted technique described in Appendix F. However, in the simple analysis described below, the location shift is ignored. A center-weighted analysis, such as that described in Appendix F, would likely result in a smaller estimate of combined standard uncertainty.

Currently, a simple “between and within” approach is used to analyze the transmittance data. The actual window transmittance present during the calibrations is assumed to be between the two mean transmittances and within their uncertainties. The transmittance estimate from the first set of measurements is called T_{W1} , with standard uncertainty u_{W1} . The results from the second set of measurement are called T_{W2} and u_{W2} . The transmittance estimates and uncertainties are derived below. The transmittance T_W at the time of the test detector calibration is assumed to lie somewhere between T_{W1} and T_{W2} , with a rectangular probability distribution, so

$$T_W = \frac{T_{W1} + T_{W2}}{2}. \quad (\text{H-1})$$

Two sources contribute uncertainty to T_W . The drift in transmittance from T_{W1} to T_{W2} contributes the standard uncertainty u_D . Assuming a rectangular probability distribution,

$$u_D = \frac{|T_{W1} - T_{W2}|}{2\sqrt{3}}. \quad (\text{H-2})$$

The uncertainty in T_{W1} and T_{W2} also contributes uncertainty T_W . Assuming the uncertainties are uncorrelated, the contribution's standard uncertainty u_M is given by

$$u_M = \sqrt{\left(\frac{u_{W1}}{2}\right)^2 + \left(\frac{u_{W2}}{2}\right)^2}. \quad (\text{H-3})$$

The combined standard uncertainty u_{TW} is

$$u_{TW} = \sqrt{u_D^2 + u_M^2} = \frac{1}{2} \sqrt{\frac{(T_{W1} - T_{W2})^2}{3} + u_{W1}^2 + u_{W2}^2}. \quad (\text{H-4})$$

The transmittance estimates T_{W1} and T_{W2} and their associated uncertainties are determined similarly. Call $T(1)$ to $T(9)$ the average transmittance measured at each of the nine locations, and $u(1)$ to $u(9)$ the standard uncertainties associated with each measurement. Several transmittance measurements are performed at each location N ; and the average measurement is used as T_{WN} , and the standard deviation of the mean of the measurements is the standard uncertainty u_{SN} .

For $N = 1$ to 2, the estimated transmittance T_{WN} for the each sample set is given by the mean of the samples:

$$T_{WN} = \frac{1}{9} \sum_{i=1}^9 T(i). \quad (\text{H-5})$$

The uncertainty in T_{WN} contains contributions from both the measured variation and the uncertainty in the measurement. The standard deviation of the nine samples is used as the standard uncertainty u_{SN} in T_{WN} that is due to variation in the transmittance,

$$u_{SN} = \sqrt{\frac{1}{8} \sum_{i=1}^9 (T(i) - T_{WN})^2}. \quad (H-6)$$

The uncertainty contribution u_{MN} due to the uncertainty in the measurements themselves is determined by using the law of propagation of uncertainty, assuming the uncertainties are uncorrelated:

$$u_{MN} = \frac{1}{9} \sqrt{\sum_{i=1}^9 u(i)^2}. \quad (H-7)$$

The combined standard uncertainty u_{WN} in T_{WN} is the quadrature sum of the two uncorrelated uncertainty components:

$$u_{WN} = \sqrt{u_{SN}^2 + u_{MN}^2} = \sqrt{\frac{1}{8} \sum_{i=1}^9 (T(i) - T_{WN})^2 + \frac{1}{81} \sum_{i=1}^9 u(i)^2}. \quad (H-8)$$

The transmittance estimate and standard uncertainty for each of the two measurements sets is determined in this way, and designated, respectively, T_{W1} and u_{W1} , and T_{W2} and u_{W2} .

Table H.1 Example measured window transmittance $T_w(x,y)$, and standard uncertainty $u_w(x,y)$.

Location (x,y) mm	Initial Measurement		Final Measurement	
	Transmittance	Standard Uncertainty	Transmittance	Standard Uncertainty
(-1,1)	0.999746	0.000020	0.999703	0.000009
(0,1)	0.999657	0.000015	0.999808	0.000004
(1,1)	0.999582	0.000013	0.999805	0.000008
(-1,0)	0.999742	0.000017	0.999754	0.000013
(0,0)	0.999686	0.000014	0.999780	0.000013
(1,0)	0.999703	0.000020	0.999823	0.000011
(-1,-1)	0.999725	0.000011	0.999768	0.000014
(0,-1)	0.999691	0.000014	0.999779	0.000006
(1,-1)	0.999736	0.000015	0.999798	0.000004

H.3 Example Calculation of Transmittance

The following example uses actual window transmittance measurement data acquired with a 633 nm wavelength laser. Table H.1 shows window transmittance data acquired before the calibrations were performed (initial measurement), and the data acquired after the calibrations were complete (final measurement). There was a shift in the nominal beam location of 1 mm in the negative-x direction, but the shift does not affect the calculations, and the location given in the table is the unshifted location. Three measurements were performed at each location; the transmittance is the mean of the measurements, and the standard deviation of the mean is the standard uncertainty. Table H.2 shows the results of the final estimation of window transmittance, its uncertainty components, and combined standard uncertainty. When expressed relative to the transmittance, the combined standard uncertainty is 0.0040 %.

Table H.2 Analysis of the example window transmittance data.

Initial Transmittance Measurement		Final Transmittance Measurement		Combined Transmittance Measurement	
T_{W1}	u_{W1}	T_{W2}	u_{W2}	T_W	u_W
0.999696	0.000052	0.999780	0.000036	0.999738	0.000040

APPENDIX I. Center-Weighted Spatial Uncertainty

Dr. Chih-Ming Wang, NIST, Statistical Engineering Division

Assume that the transmittance at (x, y) , $t(x, y)$, can be modeled approximately by a smooth curve, say a quadratic form, within a given window; that is,

$$t(x, y) = \beta_0 + \beta_1 x + \beta_2 y + \beta_3 x^2 + \beta_4 xy + \beta_5 y^2, \quad (\text{I-1})$$

where β_0, \dots, β_5 are to be estimated. Define region $R = \{(x, y); (x - x_0)^2 + (y - y_0)^2 \leq r_0^2\}$. Then the observed transmission, centered at (x_0, y_0) with beam diameter $2 r_0$, is given by

$$\hat{t}(x_0, y_0, r_0) = \iint_R w(x, y) t(x, y) dx dy, \quad (\text{I-2})$$

where $w(x, y)$ is the weight function.

We consider only the Gaussian weight. The truncated bivariate Gaussian density, centered at (x_0, y_0) having standard deviation σ and circular equal-probability contours, is given by

$$w(x, y) = \frac{1}{c} \exp\left(-\left((x - x_0)^2 + (y - y_0)^2\right) / 2\sigma^2\right). \quad (\text{I-3})$$

Since the sum of the weights must be 1, i.e.,

$$\iint_R w(x, y) dx dy = 1,$$

c is found to be $(1 - e^{-2})2\pi\sigma^2$. Also, an e^2 beam implies $w(0,0)/w(0,r) = e^2$ or $\sigma = r_0/2$.

With this Gaussian weight function, we have

$$\begin{aligned} \hat{t}(x_0, y_0, r_0) &= \frac{2}{\pi(1 - e^{-2})r_0^2} \iint_R e^{-2(x^2 + y^2)/r_0^2} t(x, y) dx dy \\ &= \beta_0 + \beta_1 x_0 + \beta_2 y_0 + \beta_3 (x_0^2 + k r_0^2) + \beta_4 x_0 y_0 + \beta_5 (y_0^2 + k r_0^2), \end{aligned} \quad (\text{I-4})$$

where $k = (1 - 3e^{-2})/(4 - 4e^{-2}) = 0.1717412$. The result was obtained by transforming the integral using polar coordinates. The result allows us to estimate β_i by linear regression with dependant variable $\hat{t}(x_0, y_0, r_0)$ and independent variables $x_0, y_0, x_0^2 + k r_0^2, x_0 y_0$, and $y_0^2 + k r_0^2$.

The data in the following example used a Gaussian beam of 2 mm in diameter. The wavelength was 1550.4 nm. The centers for the nine-measurement spatial experiment were located at $(-1, 1)$, $(-1, 0)$, $(-1, -1)$, $(0, -1)$, $(1, 1)$, $(1, 0)$, and $(1, -1)$.

The beam size used was 2 mm, i.e., $r_0 = 1$ mm. The vector of dependant variables and the matrix of independent variables for this example are

$$Y = \begin{pmatrix} 0.9994575 \\ 0.9994220 \\ 0.9992600 \\ 0.9993700 \\ 0.9993225 \\ 0.9993585 \\ 0.9993780 \\ 0.9993120 \\ 0.9993750 \end{pmatrix}, X = \begin{pmatrix} 1 & -1 & 1 & 1.17174 & -1 & 1.17174 \\ 1 & -1 & 0 & 1.17174 & 0 & 0.17174 \\ 1 & -1 & -1 & 1.17174 & 1 & 1.17174 \\ 1 & 0 & 1 & 1.17174 & 0 & 1.17174 \\ 1 & 0 & 0 & 1.17174 & 0 & 0.17174 \\ 1 & 0 & -1 & 1.17174 & 0 & 1.17174 \\ 1 & 1 & 1 & 1.17174 & 1 & 1.17174 \\ 1 & 1 & 0 & 1.17174 & 0 & 0.17174 \\ 1 & 1 & -1 & 1.17174 & -1 & 1.17174 \end{pmatrix}.$$

The parameter estimates are obtained by

$$\beta = \begin{pmatrix} \beta_0 \\ \beta_1 \\ \beta_2 \\ \beta_3 \\ \beta_4 \\ \beta_5 \end{pmatrix} = (X'X)^{-1} X'Y = \begin{pmatrix} 0.9993354 \\ -1.241667 \cdot 10^{-5} \\ 3.533333 \cdot 10^{-5} \\ 1.708333 \cdot 10^{-5} \\ -4.862500 \cdot 10^{-5} \\ 1.433333 \cdot 10^{-5} \end{pmatrix}.$$

The variance-covariance matrix of the parameters can also be calculated.

Inferences can be made based on the model just derived. For example, if we are interested in estimating the bias and the uncertainty of the transmission at (0,0) with various beam sizes, the difference in transmission between beam sizes r_0 and 1 is

$$\hat{t}(0,0,r_0) - \hat{t}(0,0,1) = k(r_0^2 - 1)(\beta_3 + \beta_5). \quad (\text{I-5})$$

The standard deviation of this difference is given by

$$u_b = k|r_0^2 - 1| \sqrt{\text{var}(\beta_3) + \text{var}(\beta_5) + 2 \text{cov}(\beta_3, \beta_5)}. \quad (\text{I-6})$$

The following table displays the bias and u_b for beam diameter $2r_0 = 1$ to 5 mm, where $\hat{\sigma}$ is the uncertainty of the transmission measurements.

Table I.1 Bias and u_b .

	Beam Diameter (mm)				
	1	2	3	4	5
bias	-0.000004	0.0	0.000007	0.000016	0.000028
u_b	$0.1288 \hat{\sigma}$	0.0	$0.2147 \hat{\sigma}$	$0.5152 \hat{\sigma}$	$0.9016 \hat{\sigma}$

We can also evaluate the uncertainty in estimating spatial transmission due to errors in locating (x_0, y_0) . The variance, assuming independence of x_0 and y_0 , is given by

$$u_p^2 = \beta_1^2 \text{var}(x_0) + \beta_2^2 \text{var}(y_0) + \beta_3^2 \text{var}(x_0^2) + \beta_4^2 \text{var}(x_0 y_0) + \beta_5^2 \text{var}(y_0^2). \quad (\text{I-7})$$

A small experiment can be conducted to estimate $\text{var}(x_0)$, \dots , $\text{var}(y_0^2)$.

APPENDIX J. Estimate of LOCR Optical Receiver Absorptance

The absorptance of the LOCR's optical receiver was measured at a wavelength of 633 nm by the device's manufacturer [13]. The receiver's absorptance at other wavelengths is estimated using measurements performed on a similar optical receiver by an outside laboratory [17]. Theoretical analysis using reflection data for the paints used on the receiver show that the device should be essentially spectrally flat over the supported wavelength range (visible to infrared), so the single absorptance measurement is commonly used. Unfortunately, the results of direct measurement of the absorptance at other wavelengths, performed and published by the outside laboratory, show that the absorptance is not spectrally flat over the wavelength region of interest to within the uncertainty of the 633 nm absorptance measurement. Imperfections in the paint, such as thick spots which could become more glossy, pooling in the corners, contamination, or voids in the paint could cause such a behavior.

J.1 LOCR Receiver Absorptance at 633 nm

The manufacturer measured the absorptance measurements of the LOCR's optical receiver at a wavelength of 633 nm. Specifically, the reflection from the receiver was measured using an integrating sphere and the absorptance was calculated from the reflectance. The receiver was illuminated with an approximately Gaussian beam, produced by a spatially filtered, single-transverse-mode helium-neon laser. Multiple reflectance measurements, called $R(N)$, were acquired by comparing the receiver's absorptance with that of a reflectance standard. The measurement has a systematic error with a standard uncertainty of approximately 2 % of the reflectance.

In each measurement, the equivalent absorptance $A(N)$ is given by $A(N) = (1 - R(N))$. The absorptance data is analyzed with a Type A evaluation, which returns the mean absorptance A_{AVE} and the standard

Table J.1 LOCR receiver absorptance at 633 nm.

N	Reflectance $R(N)$	Absorptance $A(N)$
1	0.000069	0.999931
2	0.000087	0.999913
3	0.000084	0.999916
4	0.000082	0.999918
5	0.000081	0.999919
6	0.000083	0.999917
7	0.000069	0.999931
8	0.000084	0.999916
9	0.000083	0.999917
10	0.000083	0.999917
Mean (R_{AVE} and A_{AVE})	0.000081	0.999919
Standard Deviation of the Mean (u_{TA})	0.000002	0.000002

Table J.2 Receiver absorptance uncertainty calculation.

Receiver Absorptance at 633 nm (A_{R633})	Type A Uncertainty Estimate (u_{TA})	Type B Uncertainty Estimate (u_{TB})	Combined Standard Uncertainty (u_{R633})
0.999919	0.000002	0.000002	0.000003

deviation of the mean u_A . The reflectance measurements and absorptance calculation are shown in Table J.1, where the Type A evaluation is summarized.

The standard uncertainty contributed by the systematic error, called u_B , is 2 % of the average reflection R_{AVE} , so $u_B = 0.02 \cdot R_{AVE}$. The receiver absorptance at 633 nm is called A_{R633} , and has combined standard uncertainty u_{R633} , where

$$u_{R633} = \frac{1}{A_{R633}} \sqrt{u_A^2 + u_B^2}. \quad (J-1)$$

The combined standard uncertainty calculation is summarized in Table J.2.

J.2 LOCR Receiver Absorptance at Other Wavelengths

The absorptance of the LOCR's optical receiver at wavelengths other than 633 nm is estimated using a between-and-within analysis of absorptance measurement results published by an outside laboratory [17]. The LOCR's optical receiver is identical to the measured receiver except for its scale; the LOCR's receiver is twice as large as the measured receiver; the LOCR's receiver has a diameter of approximately 10 mm, while the measured receiver has a diameter of approximately 5 mm. The receivers have otherwise the same geometry and were constructed in the same way, using the same techniques and materials, by the same manufacturer. The spectral performance of the two receivers is therefore assumed to be comparable, so the measurement results for the smaller receiver are also appropriate for the LOCR's larger receiver. The laboratories' measured absorptance of the smaller receiver at 633 nm is consistent with the manufacturer's measurement at 633 nm, so the manufacturer's measurement of the absorptance of the LOCR's receiver at 633 nm is also assumed to be accurate. The manufacturer's absorptance measurements show that the LOCR's larger receiver has a slightly higher absorptance than the smaller receiver, which is likely due to the larger receiver capturing more of the laser beam that was used in the absorptance measurement.

The outside laboratory performed their absorptance measurements in a manner similar to that used by the manufacturer to measure the absorptance at 633 nm; a spatially filtered, power stabilized, approximately Gaussian laser beam was injected into the center of the receiver, and the power reflected from the receiver was measured using an integrating sphere. The laboratory stated that their laser beams had a $1/e^2$ intensity diameter of less than 1 mm; the size of the beam used in the manufacturer's measurement is not specified in the literature, but was described in conversations as having a diameter of approximately 1 mm. The measurements are therefore assumed to be equivalent.

The outside laboratory measured their receiver's absorptance at several wavelengths from 325 to 1550 nm, in addition to the 633 nm wavelength. We assume that these measurements are typical for the receiver within the wavelength region, and therefore represent a random statistical sampling of the receiver's absorptance within this wavelength region. The measurements can therefore be used to estimate the absorptance of the LOCR's receiver within this wavelength region. These assumptions cannot be proven with the information currently available, but analysis of the paint and receiver geometry

show that the absorbance should vary smoothly throughout the wavelength region, with no resonance peaks or dips; therefore the assumption that the measurements are a valid statistical sampling of the absorbance should be valid. This analysis is considered adequate given the limited amount of absorbance information available, and is much better than assuming that the absorbance measured by the manufacturer at 633 nm is appropriate for the entire range of wavelengths. In the future, absorbance measurements may be performed on the actual LOCR receiver to verify the assumptions and reduce the uncertainty in the measurement.

In reference [17] it was noted that there was a spatial variation in the optical receiver's absorbance in that a small change in the location of the beam in the receiver caused the measured absorbance to change. They stated that this variation was a major component of the uncertainty in their absorbance measurements, and speculated that much of the measured variation in the absorption was due to spatial nonuniformity in the receiver's absorbance, or the presence of optical power in addition to that in the approximately Gaussian beam. Both of these factors probably influenced their measurements, as both have been observed in the high-accuracy calibration system. The relative aperture transmittance measurements described in Appendix F show that there is significant optical power in addition to that in the Gaussian beam, and spatial nonuniformity is evident in the evaluation of the receiver alignment uncertainty, described in Appendix I. The uncertainty that results from these imperfections likely contributed to the apparent variation in the receiver absorbance as described in the publication. The uncertainty due to these imperfections may also be double-counted in this analysis, but until the receiver absorbance measurements are reproduced using the high-accuracy calibration system where these uncertainties are assessed independently; this analysis results in an uncertainty that is at least conservative.

The outside laboratories' measurement results consist of a nominal reflectance measured at several wavelengths, called $R(\lambda)$, with associated standard uncertainty $u_R(\lambda)$. The reflection data are converted to absorbance data as described for the manufacturer's measurement, resulting in the nominal absorbance $A(\lambda)$, and its standard uncertainty $u_A(\lambda)$. The nominal absorbance over the entire wavelength range, called A_{OL} , is the mean of $A(\lambda)$. The measured variation in the absorbance is accounted for in the standard uncertainty called u_{A3} , which is the standard deviation of $A(\lambda)$; the standard deviation of mean is not used here because we are interested in the standard deviation of the absorbance over the range of wavelengths, not the uncertainty in the average A_{OL} . The uncertainty in the outside laboratories' measurements is accounted for in the uncertainty component called u_{A4} , which is the average of the uncertainties $u_A(\lambda)$.

Since the LOCR's receiver has a higher absorbance at 633 nm than the outside laboratories, the results of the outside laboratories' measurements must be scaled for use with the LOCR. The scale factor k_A is the LOCR's absorbance at 633 nm divided by the absorbance of the outside laboratories' receiver at 633 nm, so $k_A = A_{R633} / A(633 \text{ nm})$. Then the nominal absorbance of the LOCR's receiver over the entire wavelength range is $A_R = k_A \cdot A_{OL}$. Since the conversion from reflection to absorbance involves a subtraction, the uncertainty is unchanged, so the uncertainties $u_R(\lambda)$ and $u_A(\lambda)$ are equal. The combined standard uncertainty in A_{OL} , u_{AOL} , is then given by the quadrature sum of the three uncorrelated uncertainties, u_{A3} , u_{A4} , and u_{R633} . The uncertainty in A_R is $u_{AR} = k_A \cdot u_{AOL}$. The results of these calculations are given in Table J.3. When the LOCR is used at a wavelength of 633 nm, the manufacturer's absorbance measurement is used, so A_{R633} and u_{AR633} are used as A_R and u_{AR} . When the LOCR is used at any other wavelength within the range from 325 to 1550 nm, A_R and u_{AR} from Table J.3 are used instead.

Table J.3 LOCR receiver absorptance at wavelengths other than 633 nm.

Wavelength (nm)	Reflectance		Absorptance	
	$R(\lambda)$	$u_R(\lambda)$	$A(\lambda)$	$u_A(\lambda)$
325	0.000123	0.000016	0.999877	0.000016
488	0.000231	0.000030	0.999769	0.000030
515	0.000177	0.000019	0.999823	0.000019
633	0.000123	0.000013	0.999877	0.000013
1550	0.000244	0.000009	0.999756	0.000009
Standard Deviation of Absorptance, and Mean Uncertainty (u_{A3} and u_{A4})			0.000057	0.000017
Mean Receiver Absorptance and Standard Uncertainty (A_{OL} and u_{AOL})			0.999820	0.000060
LOCR Receiver Absorptance and Standard Uncertainty (A_R and u_{AR})			0.999862	0.000060

APPENDIX K. Weather Station Calibration and Uncertainty

Environmental parameters are measured using a commercial weather station [35]. To obtain accurate measurements with the weather station, its temperature and barometric pressure functions are calibrated using the following techniques. The station's relative humidity sensor is not calibrated, but its uncertainty is estimated, as described below. The standard uncertainty of the weather station measurements is combined with other uncertainties to determine the combined standard uncertainty of the measured parameter, as described in section 5.4.4. For convenience, quantization in the weather station's readings is accounted for in the calibration uncertainty. Examples of the measurements acquired during a typical calibration are shown, along with the uncertainty analysis for the example data.

The weather station is fully autonomous and records the environmental readings internally. The stored readings are downloaded to the main computer for processing and archiving as desired. Additional precision analog temperature sensors are available, which are monitored by the main computer directly, through the system's DMM. These sensors are calibrated in the same manner as the weather station's temperature sensors, except that they have no quantization error.

Currently, the weather station is calibrated as well as possible using available standards. If better standards become available, especially for pressure and humidity, the calibration uncertainties will likely be reduced. However, the uncertainty of the temperature standard (a precision mercury thermometer) is probably adequate given the size of the thermal gradients and temperature fluctuations encountered in typical optical calibrations.

K.1 Internal and External Temperature Sensor Calibration

The weather station has two temperature sensors, an internal sensor which measures the temperature on the optical source breadboard where the weather station is located, and an external sensor that measures the temperature in the detector plane. When the calibration system is not in use, the sensors are placed together, in close proximity to a precision mercury thermometer, which serves as the primary standard for the calibration of the temperature sensor. Several thermometer and weather station readings are recorded over time, and used to derive a correction factor for the weather station's indicated readings. Since the time constants of the three sensors is different, the sensors are ideally allowed to reach steady-state in a stable environment. Any error due to drift in the environment during the measurement period is randomized by acquiring the sensor data at various times throughout the day.

The correction factors for the temperature sensors are occasionally entered into the weather station's internal memory, so that the station's indicated readings are more accurate. Detector calibrations are performed continually to derive the external correction factor, which accounts for ageing and drift in the temperature sensors since the previous update. To obtain the most accurate correction possible, the external correction factors are derived from measurements acquired within a few weeks before optical calibrations are performed,. Whenever the station's internal correction factors are changed, the external calibration factor must, of course, be calculated anew.

The standard uncertainty of each temperature sensor's calibration factor has two parts: the standard uncertainty of the primary standard, called u_{T1} ; and the measurement repeatability of the sensor calibration, called u_{T2} . The combined standard uncertainty of the temperature sensor calibration, called u_{KT} , is therefore the quadrature sum of the two independent standard uncertainties, u_{T1} and u_{T2} .

An additive correction factor called k_T is derived for each temperature sensor. The correction factor is such that the actual temperature is the indicated temperature plus k_T . The correction factor is derived

using a Type A analysis of the sensor calibration data. For each sensor measurement, the difference between the standard's temperature and the station's indicated temperature is given by the amount d_T . The correction factor for the sensor, k_T , is then the mean of the offsets d_T , and the standard deviation of the mean is the measurement repeatability, u_{T2} .

Table K.1. Example temperature sensor calibration.

	Internal Temperature Sensor	External Temperature Sensor
Measured Offsets, dT ($^{\circ}\text{C}$)	-0.20	0.00
	-0.25	-0.05
	-0.10	-0.10
	-0.20	-0.10
	-0.15	-0.05
	-0.10	-0.10
	-0.10	-0.10
	-0.05	N/A
Number of Measurements, N	8	7
Mean, k_T ($^{\circ}\text{C}$)	-0.144	-0.071
Standard Deviation of the Mean, u_{T2} ($^{\circ}\text{C}$)	0.024	0.015
Mercury Thermometer Standard Uncertainty, u_{T1} ($^{\circ}\text{C}$)	0.058	0.058
Combined Standard Uncertainty, u_{kT} ($^{\circ}\text{C}$)	0.063	0.060

N/A - Not available.

The uncertainty of the temperature standard, a precision mercury thermometer, is assumed to be bounded by $\pm 0.1^{\circ}\text{C}$, the finest division on the thermometer's scale. Therefore the standard uncertainty of the primary temperature standard is $u_{T1} = 0.1^{\circ}\text{C} / \sqrt{3}$. Example calibration data and the calculated calibration factor and uncertainties are given in Table K.1 for both the internal and external temperature sensors.

The average of the temperature during the optical calibration at the weather station and in the vicinity of the external sensor is desired. Call T_I the mean internal temperature, T_E the mean external temperature, and k_I and k_E the internal and external temperature correction factors with standard uncertainties u_{kI} and u_{kE} . The correction factors and their uncertainty contribution k_I , u_{kI} , k_E , and u_{kE} are equivalent to k_T and u_{kT} from Table K.1, for the internal and external sensors. The mean temperature in the detector plane (where the external sensor is placed), called T_{DP} , is given by $T_{DP} = T_E + k_E$. Also, the mean temperature at the weather station, called T_{WS} , is given by $T_{WS} = T_I + k_I$. An example of the weather station's internal and external temperature readings, acquired during an actual calibration of optical power, are shown in Figures K.1 and K.2.




Figure K.1 Internal weather station temperature during a typical optical calibration.




Figure K.2 External weather station temperature during a typical optical calibration.

The standard uncertainty of the temperature at the weather station T_{ws} is the quadrature sum of three standard uncertainties: the standard uncertainties of the internal temperature sensor calibration u_{kl} ; the measured temperature variation, called u_{vl} ; and the standard uncertainty due to quantization in the weather station's digital output, called u_{tQ} . The measured temperature variation u_{vl} is given by the standard deviation of the measured internal temperatures; standard deviation of the mean is not used, since we are assessing the amount of variation. The digitized output, shown in Figures K.1 and K.2, is clearly quantized, with an interval of about 0.1°C . We assume that the indicated readings are rounded off so that the actual reading is within $\pm 0.05^\circ\text{C}$ of the indicated reading; therefore the standard uncertainty due to the station's quantization is $u_{tQ} = 0.05 / \sqrt{3}^\circ\text{C}$.

The detector plane temperature T_{dp} has three analogous uncertainties and one additional term. As described in section 5.4.4, the uncertainty in the detector plane temperature u_{dp} includes a term for the longitudinal temperature gradient, called u_G . This temperature gradient is much less significant since the ventilation system for the argon-ion laser was improved, but is retained for completeness. The transverse temperature gradient, evident as the periodic fluctuation in Figure K.2, is accounted for in the

Table K.2 Calculated temperatures and uncertainties for the example data.

	Internal Temperature (°C)	External Temperature (°C)
Calibration factors k_I, k_E	-0.144	-0.071
Calibration factor standard uncertainty u_{kI}, u_{kE}	0.063	0.060
Mean indicated temperature T_I, T_E	25.216	25.930
Corrected temperature T_{WS}, T_{DP}	25.072	24.859
Standard uncertainty due to temperature variation u_{VI}, u_{VE}	0.037	0.079
Output quantization standard uncertainty, u_{TQ}	0.029	0.029
Standard uncertainty due to temperature gradient u_G	0.014	0.014
Combined standard uncertainty of the corrected temperature u_{WS}, u_{DP}	0.080	0.104

measurement variation u_{VE} . The longitudinal temperature gradient is estimated as 1/10 the temperature difference between the internal and external temperature sensors. The two sensors are separated by about 1 m, and a typical test detector spans about 0.1 m, so assuming that the gradient varies linearly, the temperature gradient across the detector is bounded by: $\pm(T_{DP} - T_{WS}) / 10$. Assuming a rectangular probability distribution, the estimated standard uncertainty is

$$u_G = \frac{|T_{DP} - T_{WS}|}{10\sqrt{3}}. \quad (K-1)$$

This estimation is conservative because in the real system, most of the temperature gradient occurs across the scattered light shield located between the two sensors, which acts as a heat shield in this case. Since the uncertainties are uncorrelated, the combined standard uncertainty of the detector plane temperature measurement is the quadrature sum of the uncertainty components, so

$$u_{DP} = \sqrt{u_G^2 + u_{kE}^2 + u_{VE}^2 + u_{TQ}^2}. \quad (K-2)$$

The calculated temperatures and uncertainties for the example data are shown in Table K.2.

K.2 Barometric Pressure Sensor Calibration

The weather station's barometric pressure sensor is calibrated by comparison to nearby laboratories operated by the National Center for Atmospheric Research (NCAR). The two laboratories, the NCAR Mesa Laboratory and the NCAR Foothills Laboratory, post real-time local weather information to the internet, which include barometric pressure readings. The sites are accessible at the internet addresses www.atd.ucar.edu/cgi-bin/weather.cgi?site=ml and www.atd.ucar.edu/cgi-bin/weather.cgi?site=fl.

The two NCAR laboratories are located within a few miles of the high-accuracy calibration system, at altitudes above and below the system.

The pressure at the calibration system is interpolated from the pressure at the two nearby NCAR laboratories. The interpolation is performed assuming that the pressure changes linearly with altitude. According to the web sites, the NCAR Mesa laboratories' weather station is located at an altitude of $A_M = 1,885$ m, and the Foothills laboratories' weather station is located at an altitude of $A_F = 1,625$ m. The calibration system's weather station is located at an altitude of $A_{WS} = 1,652$ m above sea level, using the National Vertical Datum of 1929; the NCAR laboratory altitudes are assumed to be relative to the same datum, since the stated altitudes are consistent with the altitude given on topographical maps of the area, which are based on this datum. The pressure at the calibration system's weather station P_{WS} , is derived from the Foothills laboratories' pressure P_F , and the Mesa laboratories' pressure, called P_M , by the following equation:

$$P_{WS} = P_F + \frac{A_{WS} - A_F}{A_M - A_F} (P_M - P_F) \quad (K-3)$$

Deviations from eq (K-3) caused by local weather systems, sensor time constants, diurnal fluctuations, or other influences are randomized by sampling the pressures at various times throughout the day, over many days, and so are accounted for in the measurement repeatability.

The pressure readings provided by the weather station are called the indicated pressure, P_{WI} . We desire an additive calibration factor called k_p , such that the actual laboratory pressure $P_{WS} = P_{WI} + k_p$. Using a Type A analysis, the best value of k_p is the mean of $P_{WS} - P_{WI}$, for the recorded values of P_{WS} and P_{WI} , and the standard deviation of the mean is the uncertainty u_{kp} . Example pressure calibration data are given in Table K.3, along with the calculated calibration factor and uncertainty.

The uncertainty in P_{WS} that is contributed by the uncertainty in the parameters used in eq (K-3), called u_B , is determined using the law of propagation of uncertainty. Since the uncertainties are uncorrelated, the combined standard uncertainty is given by

Table K.3 Weather station barometer calibration.

	Barometric Pressure
Measured Offset, $P_{WS} - P_{WI}$ (hPa)	0.36
	0.96
	0.06
	0.56
	0.56
	1.25
	0.75
Number of Measurements, N	7
Mean Offset, k_p (hPa)	0.64
Standard Deviation of Mean, u_{kp} (hPa)	0.15

$$u_B = \sqrt{\left(\frac{\partial P_{WS}}{\partial A_{WS}}\right)^2 u_{AWS}^2 + \left(\frac{\partial P_{WS}}{\partial A_F}\right)^2 u_{AF}^2 + \left(\frac{\partial P_{WS}}{\partial A_M}\right)^2 u_{AM}^2 + \left(\frac{\partial P_{WS}}{\partial P_M}\right)^2 u_{PM}^2 + \left(\frac{\partial P_{WS}}{\partial P_F}\right)^2 u_{PF}^2} \quad (K-4)$$

To evaluate u_B , the uncertainty of the stated altitudes and NCAR pressure readings are needed. We assume a standard uncertainty for the NCAR altitudes of 2 m, so $u_{AM} = u_{AF} = 2$ m. We use a standard uncertainty of 3 m for the calibration laboratories' altitude, so $u_{AWS} = 5$ m. The NCAR web pages state that the accuracy of their barometers is 3 hPa, we assume that this is their standard uncertainty, so $u_{PM} = u_{PF} = 3$ hPa. The NCAR barometers are compensated for temperatures from 0 to 50° C, so only when the ambient temperature at both NCAR sites is within this range are the performed calibrations used in the calculation of k_p and u_{kp} ; therefore the temperature at the NCAR sites is also recorded.

In Table K.4, numerically derived partial derivatives are used to evaluate eq (K-4). The derivatives are numerically derived using nominal values for the parameters: $A_{WS} = 1652$ m, $A_F = 1625$ m, $A_M = 1885$ m, $P_F = 700$ hPa, and $P_M = 770$ hPa.

Table K.4 Uncertainty in the pressure calibration that is due to uncertainty in the parameters.

	A_{WS} (m)	A_F (m)	A_M (m)	P_M (hPa)	P_F (hPa)	u_B (hPa)
$\frac{\partial P_{WS}}{\partial x}$	0.269	-0.241	-0.028	0.896	0.896	
u_x	5	2	2	3	3	4.06

The measured variation in the pressure during the calibration also contributes uncertainty to the laboratory pressure P_{WS} . The pressures recorded by the weather station during the calibration are averaged, then the mean pressure is P_{WI} , and the standard deviation is the uncertainty due to the measured variation, called u_{PV} . Example pressure measurements acquired during an actual calibration are shown in Figure K.3. While the data shown in Figure K.3 are adequately dithered, quantization in the output is assessed for completeness. From the figure, the quantization interval is 0.1 hPa, so assuming that the readings are properly rounded off, the standard uncertainty due to quantization is $u_{PQ} = 0.05 / \sqrt{3}$ hPa.




Figure K.3 The indicated barometric pressure acquired during a typical optical calibration.

The nominal pressure during the calibration, called P , is the sum of the mean indicated pressure, P_{WI} , and the pressure calibration factor, k_p . The combined standard uncertainty in P , called u_p , is the quadrature sum of the independent uncertainties u_{kp} , u_B , u_{PV} , and u_{PQ} :

$$u_p = \sqrt{u_{kp}^2 + u_B^2 + u_{PV}^2 + u_{PQ}^2} . \quad (K-5)$$

The resulting pressure and uncertainty for the example data are given in Table K.5. The data clearly show that the uncertainty in the corrected pressure is dominated by the uncertainty in the parameters used in the pressure calibration.

Table K.5 Calculated barometric pressure and uncertainty for the example data.

	Barometric Pressure (hPa)
Calibration factor, k_p	0.64
Calibration factor standard uncertainty, u_{kp}	0.15
Indicated pressure, P_I	838.67
Corrected pressure, P	839.31
Combined standard uncertainty due to parameters, u_B	4.06
Measured pressure variation, u_{PV}	0.16
Output quantization standard uncertainty, u_{PQ}	0.03
Combined standard uncertainty of the corrected pressure, u_p	4.07

As with the system weather station's temperature sensors, the calibration factor for barometric pressure can be entered into the station so that its indicated readings are accurate. If the station's internal calibration factor is changed, the external calibration factor must be calculated anew. Calculating the correction factor continually compensates for aging and drift in the system weather station's barometer.

K.3 Relative Humidity Sensor Calibration

Currently, the weather station's relative humidity sensor is not calibrated, primarily because of the lack of an available standard for relative humidity. Typically, relative humidity sensors are calibrated by enclosing the sensor in an airtight box with known relative humidity. A known relative humidity can be generated using solutions of certain salts. The weather station's sensor could be placed in such an enclosure for calibration, or a transfer standard that was calibrated using such a method could be used. Currently, the relative humidity measurements are not considered important enough to go through such a calibration. Of importance to most detectors is whether the relative humidity is high or low, the precise relative humidity is not generally needed. The calculation of the air's index of refraction is also relatively insensitive to the humidity. Therefore, the weather station manufacturer's published specification for the accuracy of the relative humidity sensor is considered adequate.

The manufacturer states that the accuracy of the weather station's internal relative humidity sensor is $\pm 5\%$, from 10 to 90 %. We conservatively assume that this is a standard uncertainty, therefore the standard uncertainty of the humidity sensor, called u_{HS} = 5 %. There is no reason to believe that this

specification is invalid, because the published specifications for the temperature and pressure sensors seem conservative when compared to the actual performance of the sensors. The accuracy specification for both the internal and external temperature sensors is stated as $\pm 0.5^\circ \text{C}$; the calibrated sensor uncertainty is typically less than 0.1°C , so the published specification seems conservative. The accuracy of the barometric pressure sensor is stated as 1.7 hPa. Due to the inaccuracy of the pressure standard used to calibrate the sensor, it cannot be determined definitively that the sensor is actually performing better than this amount, but when compared to calibrations performed with a constant external temperature at the NCAR sites, the pressures are in significantly better agreement than implied by the published accuracy specification. This is likely because the main source of uncertainty in the NCAR pressure sensors is the external temperature. Therefore we assume that the published specification for the pressure sensor's accuracy is conservative.

Further, while the weather station is designed to internally store a calibration factor for the temperature and pressure sensors, no support for an internal humidity sensor calibration factor is provided. So it appears that the manufacturer does not believe that the sensor needs calibration to meet its published specification. Therefore we assume that the published humidity sensor specification is accurate, when the relative humidity is between 10 and 90 %. The relative humidity in the calibration laboratory is never likely to be outside this range. Therefore, we use the published accuracy specification as the standard uncertainty, and $u_{\text{HS}} = 5\%$.

Example relative humidity sensor data are shown in Figure K.4. As with the other environmental sensors, the nominal relative humidity during an optical calibration, RH , is given by the mean of the relative humidity measurements. Since the sensor is not calibrated, no correction is applied to the indicated readings. The standard uncertainty in RH , called u_{RH} , is given by the quadrature sum of three independent standard uncertainties: the sensor calibration uncertainty u_{HS} ; the measured variation, called u_{HM} ; and the standard uncertainty due to quantization in the digital output, called u_{HQ} .

The measured variation u_{HM} is given by the standard deviation of the humidity measurements. From Figure K.4, the digitized output is clearly quantized by about 1 %, therefore assuming that the values are properly rounded off, the standard uncertainty due to output quantization $u_{\text{HQ}} = 0.5 / \sqrt{3}\%$. The calculated relative humidity and uncertainty for the example data are given in Table K.6.




Figure K.4 The relative humidity measured during a typical optical calibration.

Table K.6 Calculated relative humidity and uncertainty for the example data.

	Relative Humidity (%)
Mean relative humidity, RH	31.08
Measured relative humidity variation, u_{HM}	0.27
Standard uncertainty of the sensor calibration, u_{HS}	5.00
Output quantization standard uncertainty, u_{HQ}	0.58
Combined standard uncertainty of the relative humidity, u_{RH}	5.04